arXiv:1410.3745v3 [math.PR] 3 Oct 2016

FACTOR OF IID PERCOLATION ON TREES

MUSTAZEE RAHMAN

Abstract. We study invariant percolation processes on the d-regular tree that are obtained as a factor of an iid process. We show that the density of any factor of iid site percolation process with finite clusters is asymptotically at most $(\log d)/d$ as $d \to \infty$. This bound is asymptotically optimal as it can be realized by independent sets. One implication of the result is a $(1/2)$-factor approximation gap, asymptotically in d, for estimating the density of maximal induced forests in locally tree-like d-regular graphs via factor of iid processes.

1. Introduction

Let $\mathbb{T}_d$ denote the rooted d-regular tree where a distinguished vertex $\circ$ is the root. Let χ be a finite set of 'colours'. A *factor of iid* (FIID) process on $\mathbb{T}_d$ is an invariant random process Φ taking values in $\chi^{V(\mathbb{T}_d)}$ that is defined as follows. Let $f : [0,1]^{V(\mathbb{T}_d)} \to \chi$ be a measurable function with respect to the product Borel σ-algebra of $[0,1]^{V(\mathbb{T}_d)}$. Suppose f is also invariant under root-preserving automorphisms, i.e., $f(\gamma \cdot \omega) = f(\omega)$ for every $\omega \in [0,1]^{V(\mathbb{T}_d)}$ and any $\gamma \in \mathrm{Aut}(\mathbb{T}_d)$ satisfying $\gamma(\circ) = \circ$. $\big($The automorphisms act of $[0,1]^{V(\mathbb{T}_d)}$ by $\gamma \cdot \omega(v) = \omega(\gamma^{-1}v)\big)$. A *random labelling* $\mathbf{X} = (X(v), v \in V(\mathbb{T}_d))$ of $\mathbb{T}_d$ is a process consisting of independent, identically distributed random variables $X(v)$ such that each $X(v) \sim \mathrm{Uniform}[0,1]$. Since $\mathrm{Aut}(\mathbb{T}_d)$ acts transitively on $V(\mathbb{T}_d)$, for each vertex $v \in V(\mathbb{T}_d)$ let $\gamma_{v\to\circ}$ be an automorphism such that $\gamma_{v\to\circ} v = \circ$. The process $\Phi = (\Phi(v), v \in V(\mathbb{T}_d))$ is then defined by

$$\Phi(v) = f(\gamma_{v\to\circ} \cdot \mathbf{X}) \,.$$

The value of f does not depend on the choice of $\gamma_{v\to\circ}$ because f is invariant under root-preserving automorphisms. The function f is called the *factor* associated to Φ. An FIID process $\mathbf{Y}$ is called a *percolation process* if $\mathbf{Y} \in \{0,1\}^{V(\mathbb{T}_d)}$, i.e., $\mathbf{Y}$ is $\{0,1\}$–valued at each vertex.

Informally speaking, the factor is a rule that decides the value at the root by considering the randomly labelled rooted tree as the input. Given any other vertex v the value of the process at v is determined by shifting v to the root, which naturally permutes the labels, and then applying the rule for the root. The prototypical example of FIID percolation is Bernoulli site percolation: for each $p \in [0,1]$ the factor associated to Bernoulli site percolation on $\mathbb{T}_d$ of density p is $f(\omega) = \mathbf{1}_{\{\omega(\circ)\le p\}}$.

FIID processes over $\mathbb{T}_d$, or more generally over Cayley graphs of finitely generated non-amenable groups, have been of interest in probability theory, combinatorics, computer science and ergodic theory. Recently, they have been used as randomized algorithms to construct and estimate important graph parameters such as maximal independent sets, perfect matchings, large cut sets, eigenvectors, etc. See [8, 9, 10, 12, 15, 18, 20, 21, 22] and the references therein for examples. We explain in Section 2.1 how FIID processes on $\mathbb{T}_d$ serve as a model of *local algorithms* for random d-regular graphs (or for any sequence of locally tree-like d-regular graphs).

FIID processes also arise naturally in the theory of sparse graph limits as developed by Hatami, Lovász and Szegedy [16]. The authors pose the question of what processes over large graphs can be realized as FIID

Key words and phrases. factor of iid process, induced forest, percolation, random regular graphs, local algorithms.

processes on $\mathbb{T}_d$, and make the conjecture that all reasonable optimization problems on random d-regular graphs can be approximated by FIID processes on $\mathbb{T}_d$ as the size of the graphs tend to infinity [16, Conjecture 7.13]. The conjecture was refuted for maximum independent sets by Gamarnik and Sudan [13]. They showed that maximal independent sets in random d-regular graphs cannot be approximated by FIID independent sets in $\mathbb{T}_d$ for large d. More specifically, they proved that the maximal density of FIID independent sets in $\mathbb{T}_d$ is smaller than that of random d-regular graphs by a multiplicative factor of less than 1 irrespective of the graphs' size, provided that d is sufficiently large. An asymptotically optimal gap by a factor of $1/2$, as $d \to \infty$, was established in [25]. On the other hand, the conjecture holds for perfect matchings [9, 22] and covariance structures [4].

In this paper we study FIID percolation processes and show as a consequence of the main theorem that maximal induced subgraphs with bounded size components on random d-regular graphs cannot be realized as FIID percolation processes on $\mathbb{T}_d$. Due to the local tree-like nature of random d-regular graphs, induced subgraphs with bounded components will be induced forests with high probability. On the other hand, any induced subgraph of $\mathbb{T}_d$ is a forest. Our main result (Theorem 1.1) then implies that FIID induced forests in $\mathbb{T}_d$ whose components are finite with probability one cannot model the maximal induced forests in random d-regular graphs if d is large. This is interesting because the best known lower bounds to the maximum size density of induced forests in random regular graphs is derived by way of FIID percolation processes on regular trees having finite components [17]. Since independent sets are induced forests having components of size one, our result generalizes the aforementioned results about independent sets. In order to state our results precisely we introduce some terminology.

The *density* of an FIID percolation process $\mathbf{Y}$ on $\mathbb{T}_d$ is

$$\mathrm{den}(\mathbf{Y}) = \mathbb{P}\left[\mathbf{Y}(\circ) = 1\right] = \mathbb{E}\left[f(\mathbf{X})\right] .$$

The density does not depend on the choice of the root due to invariance of $\mathbf{Y}$. The *average degree* of $\mathbf{Y}$ is the expected degree of the root, after being conditioned to be in $\mathbf{Y}$:

$$\mathrm{avdeg}(\mathbf{Y}) = \mathbb{E}\left[\#\{v \in N_{\mathbb{T}_d}(\circ) : \mathbf{Y}(v) = 1\} \mid \mathbf{Y}(\circ) = 1\right].$$

The main result of this paper states that if $\mathrm{avdeg}(\mathbf{Y}) = o(\log d)$ as $d \to \infty$ then $\mathrm{den}(\mathbf{Y}) \leq (1 + o(1))\frac{\log d}{d}$. A corollary is that when the components of the subgraph induced by $\{v \in V(\mathbb{T}_d) : \mathbf{Y}(v) = 1\}$ are finite with probability one then the density of $\mathbf{Y}$ is asymptotically at most $(1 + o(1))\frac{\log d}{d}$ as $d \to \infty$. This is because if the components are finite then the average degree of $\mathbf{Y}$ is at most 2. Our bound is optimal as $d \to \infty$ because there exists FIID independent sets in $\mathbb{T}_d$ that have density asymptotically equal to $(\log d)/d$ as $d \to \infty$ [12, 20, 26].

An important concept in the proof of the aforementioned result is a measure of correlation between the value of the root and one of its neighbours in a FIID percolation process $\mathbf{Y}$. Let $\circ'$ be a fixed neighbour of the root $\circ$. The *correlation ratio* of $\mathbf{Y}$ is defined as

$$\mathrm{corr}(\mathbf{Y}) = \frac{\mathbb{P}\left[\mathbf{Y}(\circ) = 1, \mathbf{Y}(\circ') = 1\right]}{\mathbb{P}\left[\mathbf{Y}(\circ) = 1\right] \cdot \mathbb{P}\left[\mathbf{Y}(\circ') = 1\right]} = \frac{\mathbb{P}\left[\mathbf{Y}(\circ) = 1, \mathbf{Y}(\circ') = 1\right]}{\mathrm{den}(\mathbf{Y})^2} .$$

Note that the usual correlation between $\mathbf{Y}(\circ)$ and $\mathbf{Y}(\circ')$ is $\frac{\mathbb{P}\left[\mathbf{Y}(\circ)=1,\mathbf{Y}(\circ')=1\right]-\mathrm{den}(\mathbf{Y})^2}{\mathrm{den}(\mathbf{Y})\cdot(1-\mathrm{den}(\mathbf{Y}))}$. In this paper we deal instead with the correlation ratio as defined above because it arises naturally in the course of the analysis. The correlation ratio is related to the average degree as $\mathrm{avdeg}(\mathbf{Y}_d) = d\,\mathrm{den}(\mathbf{Y})\mathrm{corr}(\mathbf{Y})$. Our main result is the following.

Theorem 1.1. *Let $\{\mathbf{Y}_d\}$ be a sequence of FIID percolation processes where $\mathbf{Y}_d$ is a percolation process on $\mathbb{T}_d$. If $\mathrm{corr}(\mathbf{Y}_d) \to 0$ as $d \to \infty$ then*

$$\limsup_{d\to\infty} \frac{\mathrm{den}(\mathbf{Y}_d)}{(\log d)/d} \leq 1\,.$$

Alternatively, if $\mathrm{avdeg}(\mathbf{Y}_d) = o(\log d)$ as $d \to \infty$ then $\limsup_{d\to\infty} \frac{\mathrm{den}(\mathbf{Y}_d)}{(\log d)/d} \leq 1$.

The result stated in terms of the average degree follows from the one stated in terms of the correlation ratio because the the aforementioned formula that connects to these two quantities. Assuming, without loss of generality, that $\mathrm{den}(\mathbf{Y}_d) \geq (\log d)/d$ we deduce that if $\mathrm{avdeg}(\mathbf{Y}_d) = o(\log d)$ then $\mathrm{corr}(\mathbf{Y}_d) \to 0$.

Corollary 1.2. *Let $\{\mathbf{Y}_d\}$ be a sequence of FIID percolation processes with $\mathbf{Y}_d$ defined on $\mathbb{T}_d$. Suppose that the connected components of the subgraph of $\mathbb{T}_d$ induced by the vertices $\{v \in V(\mathbb{T}_d) : \mathbf{Y}_d(v) = 1\}$ are finite with probability one. Then*

$$\limsup_{d\to\infty} \frac{\mathrm{den}(\mathbf{Y}_d)}{(\log d)/d} \leq 1\,.$$

Let us now explain the relevance of this result to maximum induced forests in random d-regular graphs as it relates to the conjecture of Hatami, Lovász and Szegedy. For an integer $\tau \geq 1$ define $\alpha^\tau(G)$ to be the maximum value of $|S|/|G|$ where $S \subset V(G)$ has the property that the induced subgraph $G[S]$ has components of size at most τ. Let $\mathcal{G}_{n,d}$ denote a random d-regular graph on n vertices. Bayati, Gamarnik and Tetali [1] show that $\mathbb{E}\left[\alpha^1(\mathcal{G}_{n,d})\right]$ converges as $n \to \infty$. Employing their argument it can be deduced that $\mathbb{E}\left[\alpha^\tau(\mathcal{G}_{n,d})\right]$ converges to a limit $\alpha^\tau(d)$ for every τ. It is natural to consider the quantity $\alpha^{\mathrm{finite}}(d) = \sup_\tau \alpha^\tau(d)$ and compare it to the quantity

$$\alpha^{\mathrm{finite}}_{\mathrm{FIID}}(d) = \sup\{\mathrm{den}(\mathbf{Y}) : \mathbf{Y} \text{ is a percolation process on } \mathbb{T}_d \text{ with finite components}\}.$$

Due to the locally tree-like nature of random regular graphs the quantity $\alpha^{\mathrm{finite}}(d)$ measures the size density of the largest induced forests in $\mathcal{G}_{n,d}$ as $n \to \infty$. The quantity $\alpha^{\mathrm{finite}}_{\mathrm{FIID}}(d)$ measures the largest size density of induced forests in $\mathcal{G}_{n,d}$ that are generated from FIID processes (explained further in Section 2.1). The conjecture of Hatami, Lovász and Szegedy would imply that $\alpha^{\mathrm{finite}}(d) = \alpha^{\mathrm{finite}}_{\mathrm{FIID}}(d)$ for every d. Corollary 1.2 and the existence of FIID independent sets in $\mathbb{T}_d$ of density $(1 + o(1))\frac{\log d}{d}$ implies

$$\alpha^{\mathrm{finite}}_{\mathrm{FIID}}(d) = (1 + o(1))\frac{\log d}{d} \text{ as } d \to \infty.$$

It is known that $\alpha^{\mathrm{finite}}(d) = (2 + o(1))\frac{\log d}{d}$ as $d \to \infty$ [24]. It has been actually been known much earlier that the independence ration $\alpha^1(d) = (2 + o(1))\frac{\log d}{d}$ as $d \to \infty$ [3, 11]. From these results we conclude that

$$\lim_{d\to\infty} \frac{\alpha^{\mathrm{finite}}_{\mathrm{FIID}}(d)}{\alpha^{\mathrm{finite}}(d)} = \frac{1}{2}.$$

This shows a (1/2)-factor approximation gap as $d \to \infty$ for estimating the maximum density of induced forests in random d-regular graphs by using FIID induced forest in $\mathbb{T}_d$ with finite components.

Two natural questions arise following Theorem 1.1. First, is there a bound on the density of FIID percolation processes on $\mathbb{T}_d$ in terms of the correlation ratio? Suppose $\{\mathbf{Y}_d\}$ is a sequence of FIID percolation such that $\mathrm{corr}(\mathbf{Y}_d) \to c$. What can be concluded about $\mathrm{den}(\mathbf{Y}_d)$ in terms of c? Second, what is the maximal density of an FIID percolation on the *edges* of $\mathbb{T}_d$ if it has finite clusters? The answer to the second question is due to Häggström [14] who proved that the maximal density is $2/d$.

To address the first question, note that there is no bound on the density if $c = 1$. Indeed, Bernoulli percolation on $\mathbb{T}_d$ has correlation ratio 1 and density p for any desired $p \in [0, 1]$. Also, Lyons [21] constructs FIID percolation processes on $\mathbb{T}_d$ having density arbitrarily close to $1/2$ and correlation ratio $1 - O(1/\sqrt{d})$

for large d. This construction can be utilized to get FIID percolation processes on $\mathbb{T}_d$ having density of any fixed value but having correlation ratio of order $1 - O(1/\sqrt{d})$. On the other hand, if $c < 1$ then there is a bound on the density of order $O((\log d)/d)$.

Theorem 1.3. *Let $\{\mathbf{Y}_d\}$ be a sequence of FIID percolation processes where $\mathbf{Y}_d$ is a percolation process on $\mathbb{T}_d$. Let $\Psi(c) = 1 - c + c \log c$ for $c \geq 0$ ($\Psi(0) = 1$). If $\mathrm{corr}(\mathbf{Y}_d) \leq c < 1$ for every d, or $\mathrm{corr}(\mathbf{Y}_d) \geq c > 1$ for every d then for any $\epsilon > 0$ there exists $d_0 = d_0(c, \epsilon)$ such that for $d \geq d_0$,*

$$\mathrm{den}(\mathbf{Y}_d) \leq \frac{2}{\Psi(c)d} \cdot (\log d - \log\log d + 1 + \log \Psi(c) + \epsilon) \ .$$

Also, given $0 \leq c < 1$ there exists a sequence of percolations $\{\mathbf{Z}_d\}$ with $\mathrm{corr}(\mathbf{Z}_d) = c$ and

$$\mathrm{den}(\mathbf{Z}_d) = \frac{1 - o(1)}{\sqrt{1-c}} \cdot \frac{\log d}{d} \quad \text{as } d \to \infty.$$

We are not actually aware of a sequence of FIID percolation processes whose correlation ratios satisfy $\Psi(\mathbf{Y}_d) \geq c > 1$ for every d. We also expect the bounds in Theorem 1.3 to be suboptimal but do not know how to improve upon them in a general setting.

Outline of the paper. The proofs of Theorems 1.1 and 1.3 are based on the connection between FIID processes on $\mathbb{T}_d$ and their projections onto random d-regular graphs. We will explain how to project an FIID percolation on $\mathbb{T}_d$ to a random d-regular graph. This will result in a percolation process on random d-regular graphs whose local statistics, such as the density and correlation ratio, are close to that of the original with high probability. Then we will use combinatorial arguments to bound the probability of observing a percolation process on random d-regular graphs whose statistics are close to that of the original. This leads to a fundamental non-negativity condition for an entropy-type functional that can be associated to any FIID percolation process. These results are developed in Section 2, where the main result is the derivation of the aforementioned entropy inequality in Theorem 2.2. From there we proceed to show that the entropy functional is non-negative only if the density is small with respect to the correlation ratio. This is done in Section 3 to prove Theorem 1.3 and then in Section 4 to prove Theorem 1.1.

The basic strategy of the proof of Theorem 1.1 is similar to what has been done for independent sets in [25]. Our techniques improve that result to the more general question of percolation with finite clusters. Finally, we note that the entropy functional that we come upon has been studied by Bowen [7], who developed an isomorphism invariant for free group actions on measure spaces in order to resolve questions in ergodic theory. Backhausz and Szegedy [5] have also used similar entropy arguments to show that certain processes in random d-regular graphs cannot be represented on trees.

2. Entropy inequality for FIID processes

In this section we prove an inequality that is satisfied by the entropy of the local statistics of any FIID process Φ taking values in $\chi^{V(\mathbb{T}_d)}$. The key result of this section is presented in Theorem 2.2. In order to prove this theorem we first explain how to project a FIID process $\Phi \in \chi^{V(\mathbb{T}_d)}$ onto random d-regular graphs to obtain a random χ-valued colouring of random d-regular graphs such that the local statistics of the coloring are close to the local statistics of Φ with high probability (Section 2.1). This is actually what it means to model FIID processes on random regular graphs. Following this we use combinatorial arguments to count colourings of random d-regular graphs whose statistics agree with the statistics of the projection (Section 2.2). From this counting argument we can prove Theorem 2.2.

Throughout this section the values of the degree d and the set χ are fixed. We consider asymptotic analysis in the parameter $n \to \infty$, where n is the size of the graphs to which we project Φ.

We use the configuration model [2, Chapter 2.4] as the probabilistic method to sample a random d-regular graph, $\mathcal{G}_{n,d}$, on n labelled vertices. It is sampled in the following manner. Each of the n distinct vertices emit d distinct half-edges, and we pair up these nd half-edges uniformly at random. (We assume that nd is even.) These $nd/2$ pairs of half-edges can be glued into full edges to yield a random d-regular graph. Note that the resulting graph can have loops and multiple edges. There are $(nd-1)!! = (nd-1)(nd-3)\cdots 3\cdot 1$ such graphs. We denote this set of multigraphs by $G_{n,d}$; so $\mathcal{G}_{n,d}$ is a uniform random element of $G_{n,d}$.

In order to project FIID processes on $\mathbb{T}_d$ onto $\mathcal{G}_{n,d}$ it is necessary to understand the *local structure* of $\mathcal{G}_{n,d}$. It is well known (see [19, Chapter 9.2]) that the number of cycles of length ℓ in $\mathcal{G}_{n,d}$ converges in moments to a Poisson random variable with mean $(d-1)^\ell/2\ell$. Consequently, for every constant L, the expected number of cycles in $\mathcal{G}_{n,d}$ of length at most L remains bounded as $n\to\infty$. It follows from this that $\mathcal{G}_{n,d}$ is locally tree-like in the sense that the fixed size neighbourhoods of all but an order $O(1)$ number of vertices are tress with high probability. Indeed, if the r-neighbourhood, $N_G(v,r)$, of a vertex v in a graph G is not a tree then it contains a cycle of length at most $2r$. Thus, $N_G(v,r)$ is a tree if v is not within graph distance r of any cycle in G having length at most $2r$. Since the expected number of cycles in $\mathcal{G}_{n,d}$ of length at most $2r$ is bounded in n, and the r-neighbourhood of any vertex in $\mathcal{G}_{n,d}$ contains at most d^r vertices, it follows that the expected number of vertices v in $\mathcal{G}_{n,d}$ such that $N_{\mathcal{G}_{n,d}}(v,r)$ is not a tree remains bounded in n for every fixed r. Therefore, if $\circ_n$ is a uniform random vertex of $\mathcal{G}_{n,d}$ then for every $r\geq 0$

$$\mathbb{P}\left[N_{\mathcal{G}_{n,d}}(\circ_n,r)\cong N_{\mathbb{T}_d}(\circ,r)\right] = 1-O(1/n) \quad \text{as } n\to\infty\,. \tag{2.1}$$

This is what it means for $\mathcal{G}_{n,d}$ to be locally tree-like.

2.1. **Projecting FIID processes onto random regular graphs.** Let $\mathbb{T}_{d,r}$ denote the rooted subtree $N_{\mathbb{T}_d}(\circ,r)$. We say a factor f is a *block factor* if there exists a finite r such that f maps $[0,1]^{V(\mathbb{T}_{d,r})}$ into χ, i.e., f does not depend of the labels outside of $\mathbb{T}_{d,r}$. The smallest such r is the *radius* of the factor. It is well known that given an FIID process Φ on $\mathbb{T}_d$ with factor f it is possible to find block factors f_j such that the corresponding processes Φ^j converge to Φ in the weak topology [21]. As such, for any FIID percolation process $\mathbf{Y}$ there exist approximating percolation processes $\mathbf{Y}^j$ having block factors such that $\mathrm{den}(\mathbf{Y}^j)\to\mathrm{den}(\mathbf{Y})$ and $\mathrm{corr}(\mathbf{Y}^j)\to\mathrm{corr}(\mathbf{Y})$. Henceforth, we assume that all factors associated to FIID percolation processes on $\mathbb{T}_d$ are block factors.

We now explain how to project an FIID process Φ with block factor f of radius r onto a d-regular graph $G\in G_{n,d}$. The factor f maps to a finite set χ, which we take to be $\{0,1,\ldots,q\}$. We begin with a random labelling $\mathbf{X}=(X(v),v\in V(G))$ of G. Given any vertex $v\in G$ if its r-neighbourhood is a tree then set $\Phi_G(v)=f(X(u),u\in N_G(v,r))$. This is allowed since $N_G(v,r)=\mathbb{T}_{d,r}$ by assumption. Otherwise, set $\Phi_G(v)=0$. The process Φ_G is the *projection* of Φ onto G. The projection Φ_G is also called a *local algorithm* on G; for example, if Φ is an independent set in $\mathbb{T}_d$ with probability one then Φ_G would produce random independent sets in G and be called a local algorithm for independent sets in G.

If we project Φ to $\mathcal{G}_{n,d}$ then the projection $\Phi_{\mathcal{G}_{n,d}}$ has the property that its local statistics converge to the local statistics of Φ as $n\to\infty$. More precisely, let $(\circ_n,\circ_n')$ be a uniform random directed edge of $\mathcal{G}_{n,d}$. Let $\circ$ be the root of $\mathbb{T}_d$ with let $\circ'$ be a fixed neighbour. Then for any $i,j\in\chi$, the probabilities

$$\mathbb{P}\left[\Phi_{\mathcal{G}_{n,d}}(\circ_n)=i,\Phi_{\mathcal{G}_{n,d}}(\circ_n')=j\right]\to\mathbb{P}\left[\Phi(\circ)=i,\Phi(\circ')=j\right] \quad \text{as } n\to\infty\,. \tag{2.2}$$

Consequently, $\mathbb{P}\left[\Phi_{\mathcal{G}_{n,d}}(\circ_n)=i\right]\to\mathbb{P}\left[\Phi(\circ)=i\right]$ as well. Note that $\circ_n$ is a uniform random vertex of $\mathcal{G}_{n,d}$.

To see (2.2), observe that the value of the pair $(\Phi_{\mathcal{G}_{n,d}}(\circ_n),\Phi_{\mathcal{G}_{n,d}}(\circ_n'))$ depends only on the labels on the r-neighbourhood of the edge $(\circ_n,\circ_n')$. If this is a tree then the r-neighbourhood of the edge $(\circ_n,\circ_n')$ is isomorphic to the r-neighbourhood of $(\circ,\circ')$. Therefore, $(\Phi_{\mathcal{G}_{n,d}}(\circ_n),\Phi_{\mathcal{G}_{n,d}}(\circ_n'))$ agrees with $(\Phi(\circ),\Phi(\circ'))$

if the labels on $N_{\mathcal{G}_{n,d}}((\circ_n, \circ'_n), r)$ are lifted to $N_{\mathbb{T}_d}((\circ, \circ'), r)$ in the natural way. On the other hand, if $N_{\mathcal{G}_{n,d}}((\circ_n, \circ'_n), r)$ is not a tree then $N_{\mathcal{G}_{n,d}}(\circ_n, r+1)$ is also not a tree since $N_{\mathcal{G}_{n,d}}((\circ_n, \circ'_n), r) \subset N_{\mathcal{G}_{n,d}}(\circ_n, r+1)$. As $\mathcal{G}_{n,d}$ is locally tree-like as in (2.1), we have that $\mathbb{P}\left[N_{\mathcal{G}_{n,d}}(\circ_n, r+1) \ncong \mathbb{T}_{d,r+1}\right] = O(1/n)$. As a result, it is straightforward to conclude that $\left|\mathbb{P}\left[\Phi_{\mathcal{G}_{n,d}}(\circ_n) = i, \Phi_{\mathcal{G}_{n,d}}(\circ'_n) = j\right] - \mathbb{P}\left[\Phi(\circ) = i, \Phi(\circ') = j\right]\right| = O(1/n)$.

In the proof of Theorem 2.2 we need more than the convergence of local statistics. We require that the random quantities $\#\{(u,v) \in E(\mathcal{G}_{n,d}) : \Phi_{\mathcal{G}_{n,d}}(u) = i, \Phi_{\mathcal{G}_{n,d}}(v) = j\}$ are close to their expectations. Here, (u,v) refers to a directed edge so $(u,v) \neq (v,u)$ unless $u = v$. Note that

$$\mathbb{P}\left[\Phi_{\mathcal{G}_{n,d}}(\circ_n) = i, \Phi_{\mathcal{G}_{n,d}}(\circ'_n) = j\right] = \mathbb{E}\left[\frac{\#\{(u,v) \in E(\mathcal{G}_{n,d}) : \Phi_{\mathcal{G}_{n,d}}(u) = i, \Phi_{\mathcal{G}_{n,d}}(v) = j\}}{nd}\right].$$

In other words, we require that the empirical values of the local statistics of $\Phi_{\mathcal{G}_{n,d}}$ are close to their expected value with high probability. To show this we use the Hoeffding–Azuma concentration inequality [2, Theorem 1.20]. Consider a product probability space (Ω^n, μ^n) and a bounded function $h : \Omega^n \to \mathbb{R}$ that is L-Lipschitz with respect to the Hamming distance in Ω^n. The Hoeffding-Azuma inequality states that

$$\mu^n\left(\mid h - \mathbb{E}_{\mu^n}[h] \mid > \lambda\right) \leq 2e^{-\frac{\lambda^2}{2L^2 n}}.$$

To apply this to our case, fix a graph $G \in G_{n,d}$ and consider the probability space generated by a random labelling $\mathbf{X}$ of G. The function h is taken to be $\#\{(u,v) \in E(G) : \Phi_G(u) = i, \Phi_G(v) = j\}$. We now verify that h is Lipschitz as a function of the labels with Lipschitz constant at most $2d^{r+1}$. If the value $X(w)$ of the label at w is changed to $X'(w)$ then the value of the pair $(\Phi_G(u), \Phi_G(v))$ can only change for edges (u,v) that are within graph distance r of w. Otherwise, the labels used to evaluate $(\Phi_G(u), \Phi_G(v))$ does not contain $X'(w)$. Therefore, the number of directed edges (u,v) where the value of Φ_G can change is at most $2d^{r+1}$, which is a trivial upper bound on the number of directed edges that meet $N_G(w,r)$. This implies that our function is $(2d^{r+1})$-Lipschitz.

From the Hoeffding–Azuma inequality we conclude that

$$\begin{aligned}\mathbb{P}\Bigg[\Bigg|&\frac{\#\{(u,v) \in E(\mathcal{G}_{n,d}) : \Phi_{\mathcal{G}_{n,d}}(u) = i,\ \Phi_{\mathcal{G}_{n,d}}(v) = j\}}{nd} \\ &- \mathbb{P}\left[\Phi_{\mathcal{G}_{n,d}}(\circ_n) = i, \Phi_{\mathcal{G}_{n,d}}(\circ'_n) = j\right]\Bigg| > \epsilon\Bigg] \\ &\leq 2e^{-\frac{\epsilon^2 n}{8d^{2r}}}.\end{aligned}$$

Due to convergence of the expected value of local statistics as shown in (2.2) we can replace the quantity $\mathbb{P}\left[\Phi_{\mathcal{G}_{n,d}}(\circ_n) = i, \Phi_{\mathcal{G}_{n,d}}(\circ'_n) = j\right]$ in the inequality above with the quantity $\mathbb{P}\left[\Phi(\circ) = i, \Phi(\circ') = j\right]$ at the expense of adding a term err_n such that $\mathrm{err}_n \to 0$ as $n \to \infty$. Then, after taking a union bound over all $i, j \in \chi$ we deduce from the inequality above that for some constant $C_{d,r,q}$ and error function $\mathrm{err}_n \to 0$ as $n \to \infty$, we have

$$\begin{aligned}\mathbb{P}\Bigg[\max_{i,j \in \chi}\Bigg|&\frac{\#\{(u,v) \in E(\mathcal{G}_{n,d}) : \Phi_{\mathcal{G}_{n,d}}(u) = i, \Phi_{\mathcal{G}_{n,d}}(v) = j\}}{nd} - \mathbb{P}\left[\Phi(\circ) = i, \Phi(\circ') = j\right]\Bigg| > \epsilon\Bigg] \\ &\leq e^{-\frac{\epsilon^2 n}{C_{d,r,q}}} + \mathrm{err}_n. \qquad (2.3)\end{aligned}$$

2.2. **Counting colourings of random regular graphs.** From (2.3) we see that we can prove bounds on the local statistics of Φ if we can use the structural properties of $\mathcal{G}_{n,d}$ to deduce bounds on the empirical value of the local statistics of $\Phi_{\mathcal{G}_{n,d}}$. We achieve the latter by combinatorial arguments.

On any graph $G \in G_{n,d}$ the process Φ_G creates an ordered partition Π of $V(G)$ with $q+1 = \#\chi$ cells, namely, the sets $\Pi(i) = \Phi_G^{-1}(\{i\})$ for $i \in \chi$. (It is ordered in the sense that the cells are distinguishable.) Associated to Π is its *edge profile*, defined by the quantities

$$\begin{aligned} P(i,j) &= \frac{\#\{(u,v) \in E(G) : \Phi_G(u) = i, \Phi_G(v) = j\}}{nd} \\ \pi(i) &= \frac{\#\{v \in V(G) : \Phi_G(v) = i\}}{n}. \end{aligned}$$

As we have noted earlier, if we pick a uniform random directed edge $(\circ_n, \circ'_n)$ of G then $P(i,j) = \mathbb{P}[\Phi_G(\circ_n) = i, \Phi_G(\circ'_n) = j]$ and $\pi(i) = \mathbb{P}[\Phi_G(\circ_n) = i]$. In particular, the matrix $P = \left[P(i,j)\right]_{\{i,j\in\chi\}}$ and the vector $\pi = (\pi(i))_{\{i\in\chi\}}$ are probability distributions over χ^2 and χ, respectively. Also, the marginal distribution of P along its rows is π, and P is a symmetric matrix.

Our chief combinatorial ingredient, in particular, where we use the structural property of $\mathcal{G}_{n,d}$ is in computing the expected number of ordered partitions of $\mathcal{G}_{n,d}$ that admit a given edge profile. From (2.3) we see that any FIID process Φ admits, with high probability, an ordered partition of $V(\mathcal{G}_{n,d})$ with edge profile close to $P_\Phi = \left[\mathbb{P}[\Phi(\circ) = i, \Phi(\circ') = j]\right]_{\{i,j\}}$ and $\pi_\Phi = (\mathbb{P}[\Phi(\circ) = i])_{\{i\}}$. Thus, if the probability of observing a partition of $\mathcal{G}_{n,d}$ with this edge profile is vanishingly small then the process Φ cannot exist on $\mathbb{T}_d$. Bounding the probability of observing partitions with a given edge profile is difficult. Instead, we bound the expected number of such partitions, which serves as an upper bound to the probability.

The *entropy* of a discrete probability distribution μ is

$$H(\mu) = \sum_{x \in \text{support}(\mu)} -\mu(x) \log \mu(x).$$

Let $Z(P,\pi) = Z(P,\pi,G)$ be the number of ordered partitions of $V(G)$ with edge profile (P,π).

Lemma 2.1. *For the random graph $\mathcal{G}_{n,d}$,*

$$\mathbb{E}[Z(P,\pi,\mathcal{G}_{n,d})] \le \text{poly}(n) \times e^{n\left[\frac{d}{2}H(P) - (d-1)H(\pi)\right]},$$

where poly(n) *is a polynomial factor whose degree and coefficients depend only on d and q.*

Proof. The number of candidates for ordered partitions Π on $V(\mathcal{G}_{n,d})$ that induce the edge profile (P,π) on $\mathcal{G}_{n,d}$ is equal to the number of partitions of $V(\mathcal{G}_{n,d})$ into $q+1$ distinguishable cells such that the i-th cell has size $\pi(i)n$. (Of course, we may assume that (P,π) is a valid edge profile so that the numbers $\pi(i)n$ and $P(i,j)nd$ are all integers.) The number of such partitions is given by the multinomial term $\binom{n}{\pi(i)n; i\in\chi}$.

For any such ordered partition Π, the probability that its edge profile in $\mathcal{G}_{n,d}$ agrees with (P,π) is

$$\frac{\prod_i \binom{nd\pi(i)}{ndP(i,j); j\in\chi} \prod_{\{i,j\}: i\neq j} (ndP(i,j))! \prod_i (ndP(i,i)-1)!!}{(nd-1)!!}.$$

The first product counts, for every $0 \le i \le q$, the number of ways to partition the $nd\pi(i)$ half edges from $\Pi(i)$ into distinguishable sub-cells $\Pi(i,j)$ such that $\#\Pi(i,j) = ndP(i,j)$. The half edges in $\Pi(i,j)$ are to be paired with half edges from $\Pi(j,i)$. The second and third products count the number of ways to achieve these pairings. As each configuration of pairings appear with probability $1/(nd-1)!!$, the formula above follows.

From the linearity of the expectation it follows that $\mathbb{E}[Z(P,\pi,\mathcal{G}_{n,d})]$ equals

$$\binom{n}{\pi(i)n; i\in\chi} \times \frac{\prod_i \binom{nd\pi(i)}{ndP(i,j); j\in\chi} \prod_{\{i,j\}: i\neq j} (ndP(i,j))! \prod_i (ndP(i,i)-1)!!}{(nd-1)!!}. \tag{2.4}$$

To analyze the asymptotic behaviour of (2.4) we use Stirling's approximation: $m! = \sqrt{2\pi m}(m/e)^m(1 + O(1/m))$. Note that $(m-1)!! = \frac{m!}{2^{(m/2)}(m/2)!}$ for even m. From Stirling's approximation we can easily conclude that

$$\binom{m}{\rho_i m; 0 \le i \le q} = O(m^{-q/2}) e^{mH(\rho_0,\ldots,\rho_q)} \quad \text{and} \quad \frac{m!!}{\sqrt{m!}} = \Theta(m^{-1/4})\,. \tag{2.5}$$

The $\rho_i m$ are non-negative integers such that $\rho_0 + \ldots + \rho_q = 1$. The big O and big Theta are constants that are uniform over m. From these two estimates we can simplify (2.4). The polynomial factors in n in the following come from the $m^{-q/2}$ and $m^{-1/4}$ factors in the estimates above. The first two estimates below are a result of the first estimate in (2.5) about the binomial coefficient. The last inequality comes from replacing double factorials with the second estimate provided in (2.5).

- $\binom{n}{\pi(i)n; i\in\chi} \le \text{poly}(n)\, e^{nH(\pi)}$,
- $\prod_i \binom{nd\pi(i)}{ndP(i,j); j\in\chi} \le \text{poly}(n)\, e^{nd\left[\sum_i \pi(i) H\left(\left\{\frac{P(i,j)}{\pi(i)}; 0\le j\le q\right\}\right)\right]}$,
- Finally, $\frac{\prod_{\{i,j\}: i\neq j}(ndP(i,j))! \prod_i (ndP(i,i)-1)!!}{(nd-1)!!}$ is at most

$$\text{poly}(n) \times \left[\frac{\prod_{(i,j): i\neq j}(ndP(i,j))! \prod_i (ndP(i,i))!}{(nd)!}\right]^{1/2}.$$

  Stirling's formula and algebraic simplifications imply that the right hand side of the third estimate above is bounded from above by $\text{poly}(n)\, e^{-\frac{nd}{2} H(P)}$.

The term $\sum_i \pi(i) H\big(\{P(i,j)/\pi(i);\, 0 \le j \le q\}\big)$ is the conditional entropy of P given π. Algebraic simplification shows that it equals $H(P) - H(\pi)$. Consequently,

$$\mathbb{E}\left[Z(P, \pi, \mathcal{G}_{n,d})\right] \le \text{poly}(n)\, e^{n[H(\pi) + dH(P) - dH(\pi) - (d/2)H(P)]}\,,$$

and the above simplifies to the formula in the statement of the lemma. □

We now have the ingredients to prove a key inequality about FIID processes on trees.

Theorem 2.2. *Let Φ be an FIID process on $\mathbb{T}_d$ taking values in $\chi^{V(\mathbb{T}_d)}$. Let $(\circ, \circ')$ be a directed edge of $\mathbb{T}_d$. Let P_Φ and π_Φ be the distributions of $(\Phi(\circ), \Phi(\circ'))$ and $\Phi(\circ)$, respectively. Then*

$$\frac{d}{2} H(P_\Phi) - (d-1) H(\pi_\Phi) \ge 0\,. \tag{2.6}$$

The term $\frac{d}{2} H(P_\Phi) - (d-1) H(\pi_\Phi)$ will be denoted the *entropy functional*.

Proof. For any two matrices M and M', set $||M - M'|| = \max_{i,j} |M(i,j) - M'(i,j)|$. For $0 < \epsilon < 1$ consider the event

$$A(\epsilon) = \big\{\mathcal{G}_{n,d} \text{ admits an ordered partition with edge profile } (P, \pi) \text{ satisfying } ||P - P_\Phi|| \le \epsilon\big\}.$$

Observe that there are at most $(2nd)^{(q+1)^2}$ edge profiles (P, π) that satisfy $||P - P_\Phi|| \le \epsilon$. Indeed, any such P has entries $P(i,j)$ that are rational numbers of the form a/nd with $0 \le a \le (P_\Phi(i,j) + \epsilon)nd$. Since $P_\Phi(i,j) + \epsilon \le 2$ for every entry $P_\Phi(i,j)$, there are no more than $2nd$ choices for a. As there are $(q+1)^2$ entries, the bound follows. Now taking a union bound over all edge profiles (P, π) satisfying $||P - P_\Phi|| \le \epsilon$

and using Lemma 2.1 we conclude that

$$\begin{aligned}\mathbb{P}[A(\epsilon)] &\leq \sum_{(P,\pi):\,||P-P_\Phi||\leq\epsilon} \mathbb{E}[Z(P,\pi,\mathcal{G}_{n,d})] \\ &\leq \text{poly}(n) \times \sup_{(P,\pi):\,||P-P_\Phi||\leq\epsilon} e^{n\left[\frac{d}{2}H(P)-(d-1)H(\pi)\right]}.\end{aligned}$$

The map $(P,\pi) \to (d/2)H(P) - (d-1)H(\pi)$ is continuous with respect to the norm $||\cdot||$. Continuity implies that there exists $\delta(\epsilon) \to 0$ as $\epsilon \to 0$ such that $\frac{d}{2}H(P) - (d-1)H(\pi) \leq \frac{d}{2}H(P_\Phi) - (d-1)H(\pi_\Phi) + \delta(\epsilon)$ if $||P - P_\Phi|| \leq \epsilon$. As a result,

$$\mathbb{P}[A(\epsilon)] \leq \text{poly}(n) \times e^{n[\frac{d}{2}H(P_\Phi)-(d-1)H(\pi_\Phi)+\delta(\epsilon)]}. \tag{2.7}$$

If $\frac{d}{2}H(P_\Phi) - (d-1)H(\pi_\Phi) < 0$ then by choosing ϵ small enough we can ensure that $\frac{d}{2}H(P_\Phi) - (d-1)H(\pi_\Phi) + \delta(\epsilon) < 0$. This implies from (2.7) that $\mathbb{P}[A(\epsilon)] \to 0$ as $n \to \infty$. However, this is a contradiction because $\mathbb{P}[A(\epsilon)] \to 1$ as $n \to \infty$ due to the concentration inequality in (2.3). Indeed, for any $\epsilon > 0$ we have from (2.3) that

$$\mathbb{P}[A(\epsilon)] \geq \mathbb{P}\left[||P_{\Phi_{\mathcal{G}_{n,d}}} - P_\Phi|| \leq \epsilon\right] \geq 1 - e^{-\frac{\epsilon^2 n}{C_{d,r,q}}} - \text{err}_n.$$

□

3. Proof of Theorem 1.3

We prove Theorem 1.3 before Theorem 1.1 because it serves as an introduction to the ideas used to prove the latter, which are of a more technical nature. We begin with the proof of the upper bound.

Let $\{\mathbf{Y}_d\}$ be a sequence of FIID percolation processes with $\mathbf{Y}_d$ defined on $\mathbb{T}_d$. Let $P_{\mathbf{Y}_d}$ be the 2×2 matrix whose entries are $\mathbb{P}[\mathbf{Y}_d(\circ) = i, \mathbf{Y}_d(\circ') = j]$, where $(\circ, \circ')$ is a directed edge of $\mathbb{T}_d$ and $i, j \in \{0,1\}$. This is the edge profile of $\mathbf{Y}_d$, which we can express in terms of the correlation ratio and density. For convenience we write $\rho_d = \text{corr}(\mathbf{Y}_d)$ and $\alpha_d = \text{den}(\mathbf{Y}_d)$. We have

$$P_{\mathbf{Y}_d} = \begin{bmatrix} 1 - 2\alpha_d + \rho_d\alpha_d^2 & \alpha_d - \rho_d\alpha_d^2 \\ \alpha_d - \rho_d\alpha_d^2 & \rho_d\alpha_d^2 \end{bmatrix}.$$

Set $h(x) = -x\log x$ for $0 \leq x \leq 1$ with the convention that $h(0) = 0$. The entropy functional is

$$\begin{aligned}&\frac{d}{2}H(P_{\mathbf{Y}_d}) - (d-1)H(\pi_{\mathbf{Y}_d}) = \\ &\frac{d}{2}\left[h(\rho_d\alpha_d^2) + 2h(\alpha_d - \rho_d\alpha_d^2) + h(1 - 2\alpha_d + \rho_d\alpha_d^2)\right] - (d-1)[h(\alpha_d) + h(1-\alpha_d)].\end{aligned} \tag{3.1}$$

We will first derive an upper bound to the entropy functional and then use the non-negativity of the entropy functional from Theorem 2.2 to get a bound on α_d. We will use the following properties of h.

Lemma 3.1. *The function $h(x) = -x\log x$ for $x > 0$ and $h(0) = 0$ satisfies the following.*

$$h(xy) = xh(y) + yh(x) \text{ for all } x, y \geq 0. \tag{3.2}$$

$$x - \frac{x^2 + x^3}{2} \leq h(1-x) \leq x - \frac{x^2}{2} \text{ for all } 0 \leq x \leq 1. \tag{3.3}$$

Proof. The identity in (3.2) follows from definition. The inequalities in (3.3) follow from Taylor expansion. They are valid for $x = 1$. For $0 \leq x < 1$, note that $-\log(1-x) = \sum_k x^k/k$. Hence, $-\log(1-x) \geq x + x^2/2$, which implies that $h(1-x) \geq x - (1/2)x^2 - (1/2)x^3$. Similarly, $-\log(1-x) \leq x + (1/2)x^2 + (1/3)x^3(1 + x + x^2\cdots)$, which shows that $-\log(1-x) \leq x + (1/2)x^2 + x^3/(3(1-x))$ for $0 \leq x < 1$. Consequently, $h(1-x) \leq x - (1/2)x^2 - (1/6)x^3 \leq x - (1/2)x^2$. □

Lemma 3.2. *The entropy function for $\mathbf{Y}_d$ satisfies the following upper bound.*

$$\frac{d}{2}\,H(P_{\mathbf{Y}_d}) - (d-1)H(\pi_{\mathbf{Y}_d}) \le -\frac{d}{2}\Psi(\rho_d)\alpha_d^2 + \alpha_d + h(\alpha_d) + \frac{(2\rho_d+1)d}{2}\alpha_d^3. \tag{3.4}$$

Proof. The entropy functional is the r.h.s. of (3.1), which we analyze term by term. The following inequality is derived by applying (3.2) for the equality and then using that $h(1-x) \le x$ from (3.3) with $x = \rho_d\alpha_d$ for the inequality.

$$\begin{aligned} h(\rho_d\alpha_d^2) + 2h(\alpha_d - \rho_d\alpha_d^2) &= h(\rho_d)\alpha_d^2 + 2\alpha_d h(1-\rho_d\alpha_d) + 2h(\alpha_d) \\ &\le [h(\rho_d) + 2\rho_d]\alpha_d^2 + 2h(\alpha_d). \end{aligned}$$

Similarly, the upper bound to $h(1-x)$ from (3.3) with $x = 2\alpha_d - \rho_d\alpha_d^2$ gives

$$h(1 - 2\alpha_d + \rho_d\alpha_d^2) \le 2\alpha_d - \rho_d\alpha_d^2 - \frac{1}{2}(2\alpha_d - \rho_d\alpha_d^2)^2\,.$$

Simplification then implies

$$h(\rho_d\alpha_d^2) + 2h(\alpha_d - \rho_d\alpha_d^2) + h(1 - 2\alpha_d + \rho_d\alpha_d^2) \le [h(\rho_d) + \rho_d - 2]\alpha_d^2 + 2[\alpha_d + h(\alpha_d) + \rho_d\alpha_d^3].$$

Note that $h(\rho_d) + \rho_d - 1 = -\Psi(\rho_d)$, if we recall that $\Psi(c) = c\log c - c + 1$. Also, the lower bound to $h(1-x)$ from (3.3) gives $h(1-\alpha_d) \ge \alpha_d - (1/2)\alpha_d^2 - (1/2)\alpha_d^3$. Therefore, the r.h.s. of (3.1) satisfies

$$\begin{aligned} \frac{d}{2}\left[h(\rho_d\alpha_d^2) + 2h(\alpha_d - \rho_d\alpha_d^2) + h(1 - 2\alpha_d + \rho_d\alpha_d^2)\right] - (d-1)\left[h(\alpha_d) + h(1-\alpha_d)\right] \le \\ -\frac{d}{2}\Psi(\rho_d)\alpha_d^2 + \alpha_d + h(\alpha_d) + \frac{(2\rho_d+1)d}{2}\alpha_d^3. \end{aligned}$$

□

The top order terms from the r.h.s. of (3.4) are $-(d/2)\Psi(\rho_d)\alpha_d^2$ and $h(\alpha_d)$, which are of the same order only if α_d is of order $(\log d)/d$. Writing $\alpha_d = \beta_d\frac{\log d}{d}$ the inequality (3.4) translates to

$$\frac{d}{2}\,H(P_{\mathbf{Y}_d}) - (d-1)H(\pi_{\mathbf{Y}_d}) \le \left(-\frac{\Psi(\rho_d)}{2}\beta_d + 1 + \frac{1 - \log\log d - \log\beta_d}{\log d} + \frac{(2\rho_d+1)\log d}{2d}\beta_d^2\right)\frac{\beta_d\log^2 d}{d}.$$

From the non-negativity of the entropy functional in Theorem 2.2 and the inequality above we conclude that

$$-\frac{\Psi(\rho_d)}{2}\beta_d + 1 + \frac{1 - \log\log d - \log\beta_d}{\log d} + \frac{(2\rho_d+1)\log d}{2d}\beta_d^2 \ge 0. \tag{3.5}$$

Recall that in order to prove Theorem 1.3 we must show that for any $\epsilon > 0$ there is a d_0 such that if $d \ge d_0$ then $\beta_d \le \frac{2}{\Psi(c)}\big(1 - \frac{\log\log d - 1 - \log\Psi(c) - \epsilon}{\log d}\big)$. If $\beta_d \le 1/\Psi(c)$ then this conclusion is true for all large d even for $\epsilon = 0$ so long as $c \ne 1$ (in which case $\Psi(c) = 0$). Therefore, we may assume that $\beta_d \ge 1/\Psi(c)$ for every d. If $\beta_d \ge 1/\Psi(c)$ then the inequality in (3.5) implies that

$$-\frac{\Psi(\rho_d)}{2}\beta_d + 1 + \frac{1 - \log\log d + \log\Psi(c)}{\log d} + \frac{(2\rho_d+1)\log d}{2d}\beta_d^2 \ge 0$$

because $\log\beta_d \ge -\log\Psi(c)$. Hence, we must have that $q(\beta_d) \ge 0$, where

$$q(x) = \frac{(2\rho_d+1)\log d}{2d}x^2 - \frac{\Psi(\rho_d)}{2}x + 1 + \frac{1 + \log\Psi(c) - \log\log d}{\log d}.$$

In the following we prove that under the assumptions of Theorem 1.3 β_d is at most the smaller root of the quadratic polynomial q, provided that d is sufficiently large. Then we provide an upper bound for the smaller root of q to conclude the proof of the theorem.

The two roots of q are

$$\frac{d}{(2\rho_d+1)\log d}\left(\frac{\Psi(\rho_d)}{2} \pm \sqrt{\frac{\Psi^2(\rho_d)}{4} - \frac{(4\rho_d+2)}{d}(\log d - \log\log d + \log\Psi(c) + 1)}\right). \tag{3.6}$$

We first show that q has two real roots for sufficiently large d. From Lemma 3.3 below we have that

$$\frac{\Psi(\rho_d)^2}{4\rho_d+2} \geq \frac{\Psi(c)^2}{6} \text{ if } \rho_d \leq c < 1, \text{ and } \frac{\Psi(\rho_d)^2}{4\rho_d+2} \geq \frac{\Psi(c)^2}{6c} \text{ if } \rho_d \geq c > 1.$$

Consequently, for both the cases $\rho_d \leq c < 1$ or $\rho_d \geq c > 1$ there exists a $d_1 = d_1(c)$ such that for $d \geq d_1$,

$$\frac{\Psi^2(\rho_d)}{4} - \frac{4\rho_d+2}{d}(\log d - \log\log d + \log\Psi(c) + 1) > 0 \text{ as } \frac{\log d - \log\log d + \log\Psi(c) + 1}{d} \to 0.$$

The quadratic q has two real roots when the inequality above holds because of the formula for the roots of q given in (3.6).

Assuming for $d \geq d_1$, so that q has two real roots, we now show that β_d is not larger than the smaller root of q for all large d. Lemma 3.4 below implies that $\beta_d = o(d/\log d)$ as $d \to \infty$. Therefore, given c there exists d_2 such that for $d \geq d_2$ we have $\beta_d \leq \min\{\Psi(c), \Psi(c)/c\}\frac{d}{9\log d}$. But note from (3.6) that the larger root of q is at least $\frac{d}{(2\rho_d+1)\log d} \cdot \frac{\Psi(\rho_d)}{2}$, and this quantity is at least $\min\{\Psi(c), \Psi(c)/c\}\frac{d}{6\log d}$ due to the lower bound on $\Psi(\rho_d)/(2\rho_d+1)$ provided in Lemma 3.3 below. Therefore, β_d can not be bigger than the larger root of q. Then β_d is bounded from above by the smaller root as q is negative between its two roots.

We have deduced that for $d \geq \max\{d_1, d_2\}$, β_d is bounded from above by the smaller root of q. An elementary simplification shows that the smaller root equals

$$\frac{2(1 - \frac{\log\log d}{\log d} + \frac{\log\Psi(c)}{\log d} + \frac{1}{\log d})}{\frac{\Psi(\rho_d)}{2} + \sqrt{\frac{\Psi^2(\rho_d)}{4} - \frac{(4\rho_d+2)}{d}(\log d - \log\log d + \log\Psi(c) + 1)}}.$$

Using $\sqrt{1-x} \geq 1 - x$ for $0 \leq x \leq 1$, the above may be bounded to show that for $d \geq \max\{d_1, d_2\}$,

$$\beta_d \leq \frac{2}{\Psi(\rho_d)}\left(1 + \frac{1 + \log\Psi(c) - \log\log d}{\log d} + O\big(\frac{(2\rho_d+1)\log d}{\Psi^2(\rho_d)d}\big)\right).$$

The upper bound of Theorem 1.3 now follows because the term $\frac{(2\rho_d+1)\log d}{\Psi^2(\rho_d)d}$ is of order $O(\log d/d)$ if ρ_d is uniformly bounded away from 1 in d. This follows from the lower bound on $\Psi(\rho_d)$ given in Lemma 3.3 below.

Lemma 3.3. *The function Ψ satisfies the following lower bounds.*

- $\Psi(\rho_d) \geq \Psi(c)$ *if* $\rho_d \leq c < 1$ *or if* $\rho_d \geq c > 1$.
- *If* $\rho_d \leq c < 1$ *then* $\Psi(\rho_d)/(2\rho_d+1) \geq \Psi(c)/3$.
- *If* $\rho_d \geq c > 1$ *then* $\Psi(\rho_d)/(2\rho_d+1) \geq \Psi(c)/3c$.

Proof. Differentiating $\Psi(x)$ shows that it is decreasing for $x \in [0,1]$ and increasing for $x \in [1,\infty)$. This implies the claim in the first bullet. Now, suppose that $\rho_d \leq c < 1$. Then $\Psi(\rho_d) \geq \Psi(c)$ and $2\rho_d + 1 \leq 3$, so $\Psi(\rho_d)/(2\rho_d+1) \geq \Psi(c)/3$, as required. Next, suppose that $\rho_d \geq c > 1$. Then $2\rho_d + 1 \leq 3\rho_d$ and so $\Psi(\rho_d)/(2\rho_d+1) \geq \Psi(\rho_d)/(3\rho_d)$. Observe that $\Psi(\rho_d)/\rho_d = \log\rho_d - 1 + 1/\rho_d$ and the function $x \to \log x - 1 + 1/x$ is increasing if $x \geq 1$. Hence, $\Psi(\rho_d)/(2\rho_d+1) \geq \Psi(c)/3c$ if $\rho_d \geq c$. □

Lemma 3.4. *Suppose $\rho_d = \mathrm{corr}(\mathbf{Y}_d)$ is bounded away from 1 uniformly in d, i.e., $\rho_d \notin (1-\delta, 1+\delta)$ for all d and some $\delta > 0$. Then $\alpha_d = \mathrm{den}(\mathbf{Y}_d) \to 0$ as $d \to \infty$. Equivalently, $\beta_d = o(\log d/d)$ as $d \to \infty$.*

Proof. Suppose otherwise, for the sake of a contradiction. By moving to a subsequence in d we may assume that $\alpha_d \to \alpha_\infty > 0$ and $\rho_d \to \rho_\infty \neq 1$. Then dividing (3.1) through by d and taking the limit in d we see

that

$$\frac{1}{2}\left[h(\rho_\infty\alpha_\infty^2)+2h(\alpha_\infty-\rho_\infty\alpha_\infty^2)+h(1-2\alpha_\infty+\rho_\infty\alpha_\infty^2)\right]-h(\alpha_\infty)-h(1-\alpha_\infty)\geq 0\,.$$

It is easy to show by differentiating with respect to ρ_∞ that the l.h.s. of the above is uniquely maximized at $\rho_\infty = 1$, for any fixed $\alpha_\infty > 0$. However, when $\rho_\infty = 1$, the three summands involving ρ_∞ above add up to give $2[h(\alpha_\infty)+h(1-\alpha_\infty)]$. (This conclusion, of course, follows from subadditivity of entropy and the fact that the entropy of a pair of distributions is maximal when they are independent.) Thus, when $\rho_\infty = 1$ the l.h.s. of the above inequality equals zero, but is otherwise negative for $\rho_\infty \neq 1$ and fixed $\alpha_\infty > 0$. This contradicts the inequality above and allows us to conclude that $\alpha_d \to 0$. □

3.1. **Construction of large density percolation processes with given correlation ratio.** To construct processes as in the second statement of Theorem 1.3 we interpolate between Bernoulli site percolation on $\mathbb{T}_d$ and FIID independent sets of large density. Recall that there exist FIID independent sets $\mathbf{I}_d$ on $\mathbb{T}_d$ with density $(1-o(1))\frac{\log d}{d}$ as $d\to\infty$ [20]. To construct $\mathbf{Z}_d$ from Theorem 1.3 we first fix parameters $0<p<1$ and $x>0$. Using two independent sources of random labellings of $\mathbb{T}_d$, we generate $\mathbf{I}_d$ and a Bernoulli percolation, Ber_d, having density $x\frac{\log d}{d}$. Then for each vertex v we define $\mathbf{Z}_d(v)=\mathbf{I}_d(v)$ with probability p, or $\mathbf{Z}_d(v)=\mathrm{Ber}_d(v)$ with probability $1-p$. The decisions to choose between $\mathbf{I}_d$ or Ber_d are made independently between the vertices. Then $\mathbf{Z}_d$ is an FIID percolation on $\mathbb{T}_d$. Since $\mathrm{corr}(\mathbf{I}_d)=0$ and $\mathrm{corr}(\mathrm{Ber}_d)=1$, it is easy to calculate that

$$\begin{aligned}\mathrm{den}(\mathbf{Z}_d) &= p\,\mathrm{den}(\mathbf{I}_d)+(1-p)\,\mathrm{den}(\mathrm{Ber}_d),\text{ and}\\ \mathrm{corr}(\mathbf{Z}_d) &= \frac{(1-p)^2\mathrm{den}(\mathrm{Ber}_d)^2+2p(1-p)\mathrm{den}(\mathbf{I}_d)\mathrm{den}(\mathrm{Ber}_d)}{\mathrm{den}(\mathbf{Z}_d)^2}\,.\end{aligned}$$

In the following calculation we ignore the $1-o(1)$ factor from $\mathrm{den}(\mathbf{I}_d)$ for tidiness. It does not affect the conclusion and introduces the $o(1)$ term in the statement of the lower bound for $\mathrm{den}(\mathbf{Z})$ from Theorem 1.3. Continuing from the above we see that $\mathrm{den}(\mathbf{Z}_d)=(p+(1-p)x)\frac{\log d}{d}$ and $\mathrm{corr}(\mathbf{Z}_d)=\frac{(1-p)^2x^2+2p(1-p)x}{(p+(1-p)x)^2}$. Setting the correlation ratio equal to c and solving for x in terms of p gives $x=\frac{p}{1-p}(\frac{1}{\sqrt{1-c}}-1)$. Consequently, the density of $\mathbf{Z}_d$ is $\frac{p}{\sqrt{1-c}}\cdot\frac{\log d}{d}$. By letting $p\to 1$ we deduce the required conclusion.

4. Proof of Theorem 1.1

Let $\mathbf{Y}_d$ be a sequence of FIID percolation processes with $\mathbf{Y}_d$ defined on $\mathbb{T}_d$ and the correlation ratios $\mathrm{corr}(\mathbf{Y}_d)\to 0$ as $d\to\infty$. We may conclude from Theorem 1.3 that $\limsup_{d\to\infty}\frac{\mathrm{den}(\mathbf{Y}_d)}{(\log d)/d}\leq 2$. In order to prove Theorem 1.1, where the constant 2 is replace by 1, we also use the entropy inequality in Theorem 2.2, but we apply it to many copies of a given percolation process $\mathbf{Y}_d$, where the copies are coupled in a particular way. In turns out that the entropy inequality provides better bounds when it is applied to many copies of $\mathbf{Y}_d$. This important idea is borrowed from statistical physics where it was initially observed through the *overlap structure* of the Sherrington-Kirkpatrick model [27].

To derive bounds on $\mathrm{den}(\mathbf{Y}_d)$ from the entropy inequality applied to several copies of $\mathbf{Y}_d$ we need to find a upper bound to the general entropy functional in terms of the density of each colour class (the analogue of Lemma 3.2). This bound follows from an important convexity argument. We begin with the derivation of this upper bound. Once it is established, which is presented in Lemma 4.1, the remainder of the proof follows the same argument as what has been done for independent sets in [25, Section 2.1].

4.1. **Upper bound for the entropy functional.** The subadditivity of entropy implies that for any edge profile (P,π) the entropy $H(P)\leq 2H(\pi)$, which in turn implies that $(d/2)H(P)-(d-1)H(\pi)\leq H(\pi)$.

However, this is useless towards analyzing the entropy inequality since $H(\pi) \geq 0$. In fact, this upper bound is sharp if and only if $P(i,j) = \pi(i)\pi(j)$ for every pair (i,j). However, a high density percolation process $\mathbf{Y}_d$ whose correlation ratio satisfies the hypothesis of Theorem 1.1 cannot have an edge profile with this property. In order to derive a suitable upper bound we have to bound the entropy functional subject to constraints on the edge profile induced by small correlation ratios.

We derive such a bound from convexity arguments that take into account the constraints put forth by small correlation ratios. Although our bound is not sharp, it appears to become sharp as the degree $d \to \infty$. This is the reason as to why we get an asymptotically optimal bound on the density as $d \to \infty$. When the correlation ratios do not converge to zero the bound becomes far from sharp and we cannot get bounds on the density that improve upon Theorem 1.3 in a general setting.

Before proceeding with the calculations we introduce some notation and provide the setup. Let (P, π) be an edge profile such that the entries of P are indexed by χ^2, where $\#\chi = q + 1 \geq 2$. Suppose that there is a subset $\Lambda \subset \chi^2$ (which represents the pairs (i,j) where $P(i,j)$ is very small) with the following properties.

(1) Λ is symmetric, that is, if $(i,j) \in \Lambda$ then $(j,i) \in \Lambda$.
(2) There exists an element $0 \in \chi$ such that Λ does not contain any of the pairs $(0,j)$ for every $j \in \chi$.
(3) There are positive constants $K \leq 1/(eq)$ and J such that $P(i,j) \leq K$ for every $(i,j) \in \Lambda$ and $\pi(i) \leq J$ for every $i \neq 0$ ($e = 2.78\ldots$).

For $i \in \chi$, set $\Lambda_i = \{j \in \chi : (i,j) \in \Lambda\}$ and define the quantities

$$\begin{aligned} \pi(\Lambda_i) &= \sum_{j \in \Lambda_i} \pi(j) \\ P(\Lambda_i) &= \sum_{j \in \Lambda_i} P(i,j) \\ \pi^2(\Lambda) &= \sum_i \pi(i)\pi(\Lambda_i) = \sum_{(i,j) \in \Lambda} \pi(i)\pi(j)\,. \end{aligned}$$

In our application in Section 4.3 both K and J will converge to zero as $d \to \infty$ while χ will remain fixed. So the bound stipulated on K will be satisfied for large d.

Lemma 4.1. *With the setup as above the following inequality holds for every d:*

$$\frac{d}{2}H(P) - (d-1)H(\pi) \leq H(\pi) - \frac{d}{2}\pi^2(\Lambda) + q^2\left(dK + dK\log\left(\frac{J^2}{K}\right)\right).$$

Proof. We use the concavity of the function $h(x) = -x\log x$. Using the identity $h(xy) = xh(y) + yh(x)$ from Lemma 3.1 we get

$$H(P) = \sum_{(i,j)} h\left(\pi(i)\cdot\frac{P(i,j)}{\pi(i)}\right) = \sum_{(i,j)} \pi(i) h\left(\frac{P(i,j)}{\pi(i)}\right) + H(\pi)\,.$$

To bound $\sum_i \pi(i)h(\frac{P(i,j)}{\pi(i)})$ we consider the summands for $i \in \Lambda_j$ and $i \notin \Lambda_j$ separately. Jensen's inequality applied to $h(x)$ implies that

$$\sum_{i \in \Lambda_j} \pi(i)\, h\left(\frac{P(i,j)}{\pi(i)}\right) \leq \pi(\Lambda_j)\, h\left(\frac{P(\Lambda_j)}{\pi(\Lambda_j)}\right), \text{ and} \tag{4.1}$$

$$\sum_{i \notin \Lambda_j} \pi(i)\, h\left(\frac{P(i,j)}{\pi(i)}\right) \leq (1 - \pi(\Lambda_j))\, h\left(\frac{\pi(j) - P(\Lambda_j)}{1 - \pi(\Lambda_j)}\right). \tag{4.2}$$

Since $\pi(\Lambda_0) = P(\Lambda_0) = 0$, when $j = 0$ the sum in (4.1) is empty the sum in (4.2) equals $h(\pi(0))$. We analyze the term on the r.h.s. of (4.2) when $j \neq 0$.

$$\begin{aligned}(1-\pi(\Lambda_j))\, h\left(\frac{\pi(j)-P(\Lambda_j)}{1-\pi(\Lambda_j)}\right) &= -\left(\pi(j)-P(\Lambda_j)\right)\log\left(\frac{\pi(j)-P(\Lambda_j)}{1-\pi(\Lambda_j)}\right)\\ &= h(\pi(j)-P(\Lambda_j)) + (\pi(j)-P(\Lambda_j))\log(1-\pi(\Lambda_j)).\end{aligned}$$

Notice that the term

$$h(\pi(j)-P(\Lambda_j)) = h(\pi(j)\cdot\left(1-\frac{P(\Lambda_j)}{\pi(j)})\right) = \left(1-\frac{P(\Lambda_j)}{\pi(j)}\right)h(\pi(j)) + \pi(j)\, h\left(1-\frac{P(\Lambda_j)}{\pi(j)}\right).$$

The first of the two summands on the very right, $(1-\frac{P(\Lambda_j)}{\pi(j)})\, h(\pi(j))$, equals $h(\pi(j)) + P(\Lambda_j)\log(\pi(j))$. The second, $\pi(j)h\big(1-\frac{P(\Lambda_j)}{\pi(j)}\big)$, satisfies $\pi(j)h\big(1-\frac{P(\Lambda_j)}{\pi(j)}\big) \leq P(\Lambda_j)$ because $h(1-x) \leq x$ from (3.3). Thus,

$$h\big(\pi(j)-P(\Lambda_j)\big) \leq h(\pi(j)) + P(\Lambda_j)\log(\pi(j)) + P(\Lambda_j) \text{ when } j \neq 0.$$

Now consider $(\pi(j)-P(\Lambda_j))\log(1-\pi(\Lambda_j))$. Since $\log(1-x) \leq -x$ for $0 \leq x \leq 1$, we have that $\log(1-\pi(\Lambda_j)) \leq -\pi(\Lambda_j)$. As a result,

$$(\pi(j)-P(\Lambda_j))\log(1-\pi(\Lambda_j)) \leq -\pi(j)\pi(\Lambda_j) + P(\Lambda_j).$$

We may now conclude that for $j \neq 0$,

$$(1-\pi(\Lambda_j))\, h\left(\frac{\pi(j)-P(\Lambda_j)}{1-\pi(\Lambda_j)}\right) \leq h(\pi(j)) - \pi(j)\pi(\Lambda_j) + P(\Lambda_j)\log(\pi(j)) + 2P(\Lambda_j). \tag{4.3}$$

We may also simplify the r.h.s. of (4.1): $\pi(\Lambda_j)h\big(\frac{P(\Lambda_j)}{\pi(\Lambda_j)}\big) = P(\Lambda_j)\log\big(\frac{\pi(\Lambda_j)}{P(\Lambda_j)}\big)$. Therefore, we deduce from (4.1), (4.2) and (4.3) that if $j \neq 0$ then

$$\sum_i \pi(i)h\left(\frac{P(i,j)}{\pi(i)}\right) \leq h(\pi(j)) - \pi(j)\pi(\Lambda_j) + 2P(\Lambda_j) + P(\Lambda_j)\log\left(\frac{\pi(j)\pi(\Lambda_j)}{P(\Lambda_j)}\right). \tag{4.4}$$

The last two summands on the r.h.s. of (4.4) contribute to the big O error term involving J and K, which we now demonstrate. We have that the second last summand $P(\Lambda_j) \leq qK$ by hypothesis (3). Consider the last summand in (4.4). We have that $\pi(j) \leq J$ for $j \neq 0$ by hypothesis (3). Moreover, $\pi(\Lambda_j) \leq qJ$ for $j \neq 0$ because $0 \notin \Lambda_j$ by hypothesis (2), so each $\pi(i) \leq J$ for $i \in \Lambda_j$ by hypothesis (3), and $\#\Lambda_j \leq q$. Hence, $\pi(j)\pi(\Lambda_j) \leq qJ^2$ for $j \neq 0$. The function $h(x)$ is increasing for $x \leq 1/e$. Hence, as $P(\Lambda_j) \leq qK \leq 1/e$ by hypothesis (3), we have

$$\begin{aligned}P(\Lambda_j)\log\left(\frac{\pi(j)\pi(\Lambda_j)}{P(\Lambda_j)}\right) &= h(P(\Lambda_j)) + P(\Lambda_j)\log\big(\pi(j)\pi(\Lambda_j)\big)\\ &\leq h(qK) + qK\log(qJ^2)\\ &= qK\log\left(\frac{J^2}{K}\right).\end{aligned}$$

Applying these bounds to the last two summands on the r.h.s. of (4.4) we get that that for $j \neq 0$,

$$\sum_i \pi(i)h\left(\frac{P(i,j)}{\pi(i)}\right) \leq h(\pi(j)) - \pi(j)\pi(\Lambda_j) + 2qK + qK\log\left(\frac{J^2}{K}\right).$$

Summing over all j now gives

$$\begin{aligned}\sum_{(i,j)} \pi(i) h\left(\frac{P(i,j)}{\pi(i)}\right) &\leq h(\pi(0)) + \sum_{j\neq 0} [h(\pi(j)) - \pi(j)\pi(\Lambda_j)] + 2q^2\left(K + K\log\left(\frac{J^2}{K}\right)\right) \\ &= H(\pi) - \pi^2(\Lambda) + 2q^2\left(K + K\log\left(\frac{J^2}{K}\right)\right).\end{aligned}$$

Consequently, $H(P) \leq 2H(\pi) - \pi^2(\Lambda) + 2q^2\big(K + K\log(J^2/K)\big)$, and this implies the lemma. □

4.2. **Tools and setup of the proof.** In this section we set up important tools and terminology that we will use to prove Theorem 1.1 in the following section.

4.2.1. *Many overlapping processes from a given percolation process.* To prove Theorem 1.1 we construct a coupled sequence of exchangeable percolation processes $\mathbf{Y}_d^0, \mathbf{Y}_d^1, \ldots$ from a given percolation process $\mathbf{Y}_d$. We then apply the entropy inequality to the FIID process $(\mathbf{Y}_d^1, \ldots, \mathbf{Y}_d^k)$ to derive bounds on $\mathrm{den}(\mathbf{Y}_d)$ that improve upon Theorem 1.3 by taking k to be arbitrarily large. The process $\mathbf{Y}_d^i$ is defined by varying the random labels that serve as input to the factor of $\mathbf{Y}_d$. The labels are varied by first generating a random subset of vertices of $\mathbb{T}_d$ via a Bernoulli percolation and then repeatedly re-randomizing the labels over this subset to produce the random labelling for the $\mathbf{Y}_d^i$. The processes are coupled since their inputs agree on the complement of this random subset. The details follow.

Let $\mathbf{X}_d^0, \mathbf{X}_d^1, \mathbf{X}_d^2, \ldots$ be independent random labellings of $\mathbb{T}_d$. Fix a parameter $0 \leq p \leq 1$, and let Ber_d^p be a Bernoulli percolation on $\mathbb{T}_d$ having density p. If $f_{\mathbf{Y}_d}$ is the factor associated to $\mathbf{Y}_d$ then the FIID percolation process $\mathbf{Y}_d^i$ has factor $f_{\mathbf{Y}_d}$ but the random labelling used for it, say $\mathbf{W}^i = (W^i(v), v \in V(\mathbb{T}_d))$, is defined by $W^i(v) = X_d^i(v)$ if $\mathrm{Ber}_d^p(v) = 1$ and $W^i(v) = X^0(v)$ if $\mathrm{Ber}_d^p(v) = 0$.

For $k \geq 1$ consider the FIID process $(\mathbf{Y}_d^1, \ldots, \mathbf{Y}_d^k)$, which takes values in $\chi^{V(\mathbb{T}_d)}$ where $\chi = \{0,1\}^k$. By identifying elements of $\{0,1\}^k$ with subsets of $\{1, \ldots, k\}$ the edge profile of this process can be described as follows. Let $(\circ, \circ')$ be a fixed edge of $\mathbb{T}_d$. For subsets $S, T \subset \{1, \ldots, k\}$,

$$\begin{aligned}\pi(S) &= \mathbb{P}\left[\mathbf{Y}_d^i(\circ) = 1 \text{ for } i \in S \text{ and } \mathbf{Y}_d^i(\circ) = 0 \text{ for } i \notin S\right] \\ P(S,T) &= \mathbb{P}\Big[\mathbf{Y}_d^i(\circ) = 1, \mathbf{Y}_d^j(\circ') = 1 \text{ for } i \in S, j \in T \\ &\qquad \text{and } \mathbf{Y}_d^i(\circ) = 0, \mathbf{Y}_d^j(\circ') = 0 \text{ for } i \notin S, j \notin T\Big]\end{aligned}$$

Observe that $\pi(\{i\}) \leq \mathrm{den}(\mathbf{Y}_d)$ for every i, and if $S \cap T \neq \emptyset$ then for any fixed $i \in S \cap T$ we have

$$P(S,T) \leq \mathbb{P}\left[\mathbf{Y}_d^i(\circ) = 1, \mathbf{Y}_d^i(\circ') = 1\right] \leq \mathrm{corr}(\mathbf{Y}_d)\,\mathrm{den}(\mathbf{Y}_d)^2. \tag{4.5}$$

We describe some important properties of the coupled processes that will be used to complete the proof of Theorem 1.1. Since the statement of the theorem is about the scaled density $\pi(\{1\})\frac{d}{\log d}$, we find a probabilistic interpretation for these scaled quantities. We use the coupling to find such an interpretation for the ratio $\mathbb{P}\left[\mathbf{Y}_d^i(\circ) = 1 \text{ for every } i \in S\right]/\mathrm{den}(\mathbf{Y}_d)$. Define the normalized density $\alpha(S)$ for any non empty and finite set $S \subset \{1, 2, 3, \ldots\}$ by

$$\alpha(S)\frac{\log d}{d} = \mathbb{P}\left[\mathbf{Y}_d^i(\circ) = 1 \text{ for every } i \in S\right]. \tag{4.6}$$

Since the coupled percolation processes are exchangeable, $\alpha(S) = \alpha(\{1, \ldots, \#S\})$ for all S. For convenience, we write $\alpha_{i,d,p} = \alpha(\{1, \ldots, i\})$ and call these the *intersection densities* of the processes $\mathbf{Y}_d^1, \mathbf{Y}_d^2, \ldots$ (although, these densities are normalized by the factor of $(\log d)/d$). When we apply the entropy inequality to $(\mathbf{Y}_d^1, \ldots, \mathbf{Y}_d^k)$ we will get an expression in terms of the intersection densities of the k percolation processes.

In order to analyze this expression we have to realize the ratios $\alpha_{i,d,p}/\alpha_{1,d}$ as the moments of a random variable $Q_{d,p}$. We define $Q_{d,p}$ in the following.

4.2.2. *The stability variable.* The random variable $Q_{d,p}$ is defined on a new probability space, which is obtained from the original probability space generated by the random labels $\mathbf{X}_d$ by essentially restricting to the support of the factor $f_{\mathbf{Y}^d}$. Formally, the new sample space is the set $\{\mathbf{Y}_d^0(\circ) \equiv 1\}$ considered as a subset of the joint sample space of $\mathbf{X}_d^0, \mathbf{X}_d^1, \ldots$, and Ber_d^p. The new σ-algebra is the restriction of the σ-algebra generated by $\mathbf{X}_d^0, \mathbf{X}_d^1, \ldots$, and Ber_d^p to $\{\mathbf{Y}_d^0(\circ) \equiv 1\}$. The new expectation operator $\mathbb{E}^*$ is defined by

$$\mathbb{E}^*[U] = \frac{\mathbb{E}\left[\mathbf{Y}_d^0(\circ)\, U\right]}{\mathbb{E}\left[\mathbf{Y}_d^0(\circ)\right]}$$

for any measurable random variable U defined on $\{\mathbf{Y}_d^0(\circ) \equiv 1\}$. In the following, we write $\mathbf{Y}^i$ to stand for $\mathbf{Y}_d^i$.

Lemma 4.2. *If $\mathcal{F}$ is a σ-algebra containing the σ-algebra generated by $\mathbf{Y}^0(\circ)$, then for any random variable U defined on the original probability space we have*

$$\mathbb{E}^*[U \mid \mathcal{F}] = \mathbb{E}[U \mid \mathcal{F}]\ .$$

This is to be interpreted by restricting $\mathcal{F}$ to $\{\mathbf{Y}^0(\circ) \equiv 1\}$ on the left and the random variable $\mathbb{E}[U \mid \mathcal{F}]$ to $\{\mathbf{Y}^0(\circ) \equiv 1\}$ on the right.

Proof. Suppose that Z is a $\mathcal{F}$-measurable random variable. Then,

$$\mathbb{E}^*[Z\,\mathbb{E}[U \mid \mathcal{F}]] = \frac{\mathbb{E}\left[\mathbf{Y}^0(\circ) Z\,\mathbb{E}[U \mid \mathcal{F}]\right]}{\mathbb{E}\left[\mathbf{Y}^0(\circ)\right]} = \frac{\mathbb{E}\left[\mathbb{E}\left[\mathbf{Y}^0(\circ) Z\, U \mid \mathcal{F}\right]\right]}{\mathbb{E}\left[\mathbf{Y}^0(\circ)\right]}.$$

The last equality is because $\mathbf{Y}^0(\circ)$ and Z are $\mathcal{F}$-measurable. From the definition of conditional expectation,

$$\frac{\mathbb{E}\left[\mathbb{E}\left[\mathbf{Y}^0(\circ) Z\, U \mid \mathcal{F}\right]\right]}{\mathbb{E}\left[\mathbf{Y}^0(\circ)\right]} = \frac{\mathbb{E}\left[\mathbf{Y}^0(\circ) ZU\right]}{\mathbb{E}\left[\mathbf{Y}^0(\circ)\right]} = \mathbb{E}^*[ZU]\,.$$

The lemma follows from the definition of conditional expectation. □

Define a $[0,1]$-valued random variable $Q_{d,p} = Q_d(\mathrm{Ber}_d^p, \mathbf{X}_d^0)$, which we denote the *stability*, on the restricted probability space as follows. Set

$$Q_{d,p} = \mathbb{E}^*\left[\mathbf{Y}^1(\circ) \mid \mathbf{X}_d^0, \mathrm{Ber}_d^p\right] = \mathbb{E}\left[\mathbf{Y}^1(\circ) \mid \mathbf{X}_d^0, \mathrm{Ber}_d^p\right]\ . \tag{4.7}$$

The second equality follows from Lemma 4.2. In an intuitive sense, the stability is the conditional probability, given the root is included in the percolation process, that it remains to be included after re-randomizing the labels. We now observe that the ratio of the intersection densities are moments of the stability.

Lemma 4.3. *For every $i \geq 1$ we have that*

$$\mathbb{E}^*\left[Q_{d,p}^{i-1}\right] = \frac{\alpha_{i,d}}{\alpha_{1,d}}.$$

Proof. Note that $Q_{d,p} = \mathbb{E}\left[\mathbf{Y}^j(\circ) \mid \mathbf{X}_d^0, \mathrm{Ber}_d^p\right]$ for every j from (4.7) and symmtery. Hence,

$$\begin{aligned}\mathbb{E}^*\left[Q_{d,p}^{i-1}\right] &= \frac{\mathbb{E}\left[\mathbf{Y}^0(\circ)\left(\mathbb{E}\left[\mathbf{Y}^1(\circ) \mid \mathbf{X}_d^0, \mathrm{Ber}_d^p\right]\right)^{i-1}\right]}{\mathbb{E}\left[\mathbf{Y}^0(\circ)\right]} \\ &= \frac{\mathbb{E}\left[\mathbf{Y}^0(\circ)\,\left(\prod_{j=1}^{i-1}\mathbb{E}\left[\mathbf{Y}^j(\circ) \mid \mathbf{X}_d^0, \mathrm{Ber}_d^p\right]\right)\right]}{\mathbb{E}\left[\mathbf{Y}^0(\circ)\right]}\,.\end{aligned}$$

Observe that the $\mathbf{Y}^j$ are independent of each other conditioned on $(\mathbf{X}_d^0, \mathrm{Ber}_d^p)$. Hence,

$$\prod_{j=1}^{i-1} \mathbb{E}\left[\mathbf{Y}^j(\circ) \mid \mathbf{X}_d^0, \mathrm{Ber}_d^p\right] = \mathbb{E}\left[\prod_{j=1}^{i-1} \mathbf{Y}^j(\circ) \mid \mathbf{X}_d^0, \mathrm{Ber}_d^p\right].$$

Since $\mathbf{Y}^0$ is measurable with respect to $\mathbf{X}_d^0$ we see that

$$\begin{aligned}\mathbb{E}\left[\mathbf{Y}^0(\circ)\left(\prod_{j=1}^{i-1} \mathbb{E}\left[\mathbf{Y}^j(\circ) \mid \mathbf{X}_d^0, \mathrm{Ber}_d^p\right]\right)\right] &= \mathbb{E}\left[\mathbf{Y}^0(\circ)\prod_{j=1}^{i-1} \mathbf{Y}^j(\circ)\right] \\ &= \mathbb{P}\left[\mathbf{Y}^j(\circ) = 1,\ 0 \le j \le i-1\right].\end{aligned}$$

Thus, $\mathbb{E}^*\left[Q_{d,p}^{i-1}\right] = \frac{\mathbb{P}\left[\mathbf{Y}^j(\circ)=1,\ 0\le j\le i-1\right]}{\mathrm{den}(\mathbf{Y}_d)} = \frac{\alpha_{i,d}}{\alpha_{1,d}}$. □

We will require the following continuity lemma about the stability.

Lemma 4.4. *For each $u \ge 0$, the moment $\mathbb{E}^*\left[Q_{d,p}^u\right]$ is a continuous function of p. When $p = 0$, $\mathbb{E}^*\left[Q_{d,p}^u\right] = 1$, and when $p = 1$, $\mathbb{E}^*\left[Q_{d,p}^u\right] = \mathrm{den}(\mathbf{Y}_d)^u$.*

Proof. Recall we had assumed that the factor $f_{\mathbf{Y}_d}$ is a block factor with some radius $r = r_d < \infty$. The parameter p enters into $\mathbb{E}^*\left[Q_{d,p}^u\right]$ through the random finite set $\{v \in V(\mathbb{T}_{d,r}) : \mathrm{Ber}_d^p(v) = 1\}$. If $S \subset V(\mathbb{T}_{d,r})$ then

$$\mathbb{P}\left[\mathrm{Ber}_d^p(v) = 1 \text{ for } v \in S \text{ and } \mathrm{Ber}_d^p(v) = 0 \text{ for } v \in V(T_{d,r}) \setminus S\right] = p^{\#S}(1-p)^{\#(V(T_{d,r})\setminus S)}.$$

This is a polynomial in p, and by conditioning on the output of Ber_d^p restricted to $V(\mathbb{T}_{d,r})$, it follows that $\mathbb{E}^*\left[Q_{d,p}^u\right]$ can be expressed as a convex combination of terms that are free of p with coefficients given by these probabilities. Thus, $\mathbb{E}^*\left[Q_{d,p}^u\right]$ is a polynomial in p as well.

When $p = 0$ the process $\mathrm{Ber}_d^p \equiv 0$, and $\mathbf{Y}^0 = \mathbf{Y}^1$. Therefore, conditioning on $\mathbf{X}_d^0$ and restricting to $\{\mathbf{Y}_d^0(\circ) \equiv 1\}$ makes $Q_{d,0} \equiv 1$. When $p = 1$ the process $\mathrm{Ber}_d^p \equiv 1$. So $\mathbf{Y}^1$ becomes independent of $\mathbf{X}_d^0$. Then the conditioning has no effect and $Q_{d,1} = \mathbb{E}\left[\mathbf{Y}^1(\circ)\right] = \mathrm{den}(\mathbf{Y}_d)$. □

4.3. **Completing the proof.** Let $\{\mathbf{Y}_d\}$ be a sequence of FIID percolation processes with $\mathbf{Y}_d \in \{0,1\}^{V(\mathbb{T}_d)}$. For the sake of a contradiction we assume that $\mathrm{corr}(\mathbf{Y}_d) \to 0$ while $\mathrm{den}(\mathbf{Y}_d) \ge \alpha\frac{\log d}{d}$ for some $\alpha > 1$. (Technically speaking, we need to move to a subsequence in d, but we can assume WLOG that this is the case.) Note that it follows from Theorem 1.3 that $\alpha \le 2$. Because of this we may also assume WLOG that $\mathrm{den}(\mathbf{Y}_d) \le 10\frac{\log d}{d}$ for every d.

For fixed values of $k \ge 1$ and $0 \le p \le 1$ consider the coupling $(\mathbf{Y}_d^1, \ldots, \mathbf{Y}_d^k)$ described in Section 4.2. Denote the edge profile of this coupled process as (P_d, π_d) where $P_d = [P_d(S,T)]_{S,T\subset\{1,\ldots,k\}}$. The arguments in this section will involve asymptotics for large d with the values of k and p fixed. As we assumed that $\mathrm{den}(\mathbf{Y}_d) \le 10\frac{\log d}{d}$, for any subset $S \ne \emptyset$, the quantity

$$\pi_d(S) \le \pi(\{1\}) \le \mathrm{den}(\mathbf{Y}_d) \le 10\frac{\log d}{d}. \tag{4.8}$$

Let us write $\mathrm{corr}(\mathbf{Y}_d) = \epsilon_d$ so that $\epsilon_d \to 0$ as $d \to \infty$ by assumption. Then we have from (4.5) that

$$P_d(S,T) \le \epsilon_d \frac{\log^2 d}{d^2} \text{ for every pair } S, T \text{ such that } S \cap T \ne \emptyset. \tag{4.9}$$

We apply the upper bound from Lemma 4.1 to the entropy functional associated to the edge profile (P_d, π_d). We take $\chi = \{0,1\}^k$, the role of the element 0 is taken by the empty set, and $\Lambda = \{(S,T) : S \cap T \ne \emptyset\}$. We may take $K = 100\epsilon_d\frac{\log^2 d}{d^2}$, which follows from (4.9). Also, we may take $J = 10\frac{\log d}{d}$. Applying Lemma 4.1

we conclude that for some $d_0 = d_0(k)$, if $d \geq d_0$ then

$$\frac{d}{2}H(P_d) - (d-1)H(\pi_d) \leq H(\pi_d) - \frac{d}{2}\pi_d^2(\Lambda) + 4^k\left[\epsilon_d + h(\epsilon_d)\right]\frac{\log^2 d}{d}. \tag{4.10}$$

We bound the r.h.s. of (4.10). For $S \neq \emptyset$, define $\beta(S)$ by writing $\pi_d(S) = \beta(S)\frac{\log d}{d}$. We ignore explicitly writing the dependence of $\beta(S)$ on d. With this notation we have

$$\pi_d^2(\Lambda) = \sum_{(S,T):S\cap T\neq\emptyset} \beta(S)\beta(T)\frac{\log^2 d}{d^2}.$$

Lemma 4.5. *The entropy functional for (P_d, π_d) satisfies the following for $d \geq d_0$:*

$$\frac{d}{2}H(P_d) - (d-1)H(\pi_d) \leq \left(\sum_{S\neq\emptyset}\beta(S) - \frac{1}{2}\sum_{(S,T):S\cap T\neq\emptyset}\beta(S)\beta(T)\right)\frac{\log^2 d}{d} + 10\cdot 2^k\cdot\frac{\log d}{d} + 4^k\left[\epsilon_d + h(\epsilon_d)\right]\frac{\log^2 d}{d}.$$

Proof. We consider the r.h.s. of (4.10). For $S \neq \emptyset$, from the identity (3.2) we have that $h(\pi_d(S)) \leq \beta(S)\frac{\log^2 d}{d} + h(\beta(S))\frac{\log d}{d}$. Now, $h(\beta(S)) \leq 1$ since $h(x) \leq 1$ for all $x \geq 0$. Thus, $h(\pi_d(S)) \leq \beta(S)\frac{\log^2 d}{d} + \frac{\log d}{d}$. On the other hand,

$$\pi_d(\emptyset) = 1 - \sum_{S\neq\emptyset}\beta(S)\frac{\log d}{d}, \text{ and } \sum_{S\neq\emptyset}\beta(S)\frac{\log d}{d} \leq 10\cdot 2^k\cdot\frac{\log d}{d} \text{ since } \beta(S) \leq 10 \text{ by (4.8).}$$

Hence, $h(\pi_d(\emptyset)) \leq 10\cdot 2^k\cdot\frac{\log d}{d}$ because $h(1-x) \leq x$. Therefore,

$$H(\pi_d) - \frac{d}{2}\pi_d^2(\Lambda) \leq \left(\sum_{S\neq\emptyset}\beta(S) - \frac{1}{2}\sum_{(S,T):S\cap T\neq\emptyset}\beta(S)\beta(T)\right)\frac{\log^2 d}{d} + \frac{11\cdot 2^k\cdot\log d}{d}.$$

The desired conclusion follows from the inequality in (4.10). □

The bound provided in Lemma 4.5 depends on the coupling parameter p, which so far has been held fixed, only through the $\beta(S)$'s. We may take an infimum over $p \in [0,1]$ in the inequality given in Lemma 4.5. Moreover, the term $\epsilon_d + h(\epsilon_d) \to 0$ as $d \to \infty$. Hence, the top order term on the r.h.s. of the inequality presented in Lemma 4.5 is the term of order $\frac{\log^2 d}{d}$. Then, due to the non-negativity of the entropy functional (Theorem 2.2) we deduce from Lemma 4.5 that for every $k \geq 1$,

$$\liminf_{d\to\infty}\inf_{p\in[0,1]}\sum_{S\neq\emptyset}\beta(S) - \frac{1}{2}\sum_{(S,T):S\cap T\neq\emptyset}\beta(S)\beta(T) \geq 0. \tag{4.11}$$

For the analysis of (4.11) it is convenient to parametrize $\beta(S)$ through the intersection densities of the coupled process, as defined in (4.6), and then express (4.11) in terms of those densities. The principle of inclusion and exclusion provides the following relation between the $\beta(S)$ and the intersection densities $\alpha(S)$ defined in (4.6).

$$\alpha(S) = \sum_{T:S\subset T}\beta(T) \tag{4.12}$$

$$\beta(S) = \sum_{T:S\subset T}(-1)^{\#(T\setminus S)}\alpha(T)\,. \tag{4.13}$$

We now show that the quantity in (4.11) can be rewritten as follows.

Lemma 4.6.

$$\sum_{S\neq\emptyset}\beta(S)-\frac{1}{2}\sum_{(S,T):S\cap T\neq\emptyset}\beta(S)\beta(T)=\sum_{i}(-1)^{i-1}\binom{k}{i}(\alpha_{i,d,p}-\frac{1}{2}\alpha_{i,d,p}^2)\,. \tag{4.14}$$

Proof. The term $\sum_{S\neq\emptyset}\beta(S)$ equals $\sum_i(-1)^{i-1}\binom{k}{i}\alpha_{i,d,p}$ because both these terms are equal to the normalized density $\mathbb{P}\left[\mathbf{Y}_d^i(\circ)=1 \text{ for some } i\leq k\right]\cdot\frac{d}{\log d}$. The relations (4.12) and (4.13) imply that the quantity $\alpha(S)^2=\sum_{(T_1,T_2):S\subset T_1\cap T_2}\beta(T_1)\beta(T_2)$. Thus,

$$\begin{aligned}
\sum_{i=1}^{k}(-1)^{i-1}\binom{k}{i}\alpha_{i,d,p}^2 &= -\sum_{S\neq\emptyset}(-1)^{\#S}\alpha(S)^2\\
&= -\sum_{S\neq\emptyset}(-1)^{\#S}\sum_{(T_1,T_2):S\subset T_1\cap T_2}\beta(T_1)\beta(T_2)\\
&= -\sum_{(T_1,T_2)}\beta(T_1)\beta(T_2)\sum_{S:S\subset T_1\cap T_2,S\neq\emptyset}(-1)^{\#S}\\
&= -\sum_{(T_1,T_2):T_1\cap T_2\neq\emptyset}\beta(T_1)\beta(T_2)\sum_{i=1}^{\#T_1\cap T_2}(-1)^i\binom{\#T_1\cap T_2}{i}\,.
\end{aligned}$$

Recall the binomial identity $\sum_{i=1}^{t}(-1)^i\binom{t}{i}=(1-1)^t-1=-1$ for any integer $t\geq 1$. Using this to simplify the last term from the equations above gives

$$\sum_{i=1}^{k}(-1)^{i-1}\binom{k}{i}\alpha_{i,d,p}^2=\sum_{(S,T):S\cap T\neq\emptyset}\beta(S)\beta(T)\,,\text{ as required.}$$

□

We now express (4.11) in terms of the stability via Lemma 4.6. Let $Q_{d,p}$ denote the stability of $\mathbf{Y}_d$ as defined in (4.7). Recall that $\alpha_{i,d,p}=\alpha_{1,d}\mathbb{E}^*\left[Q_{d,p}^{i-1}\right]$ for all $i\geq 1$ from Lemma 4.3. Henceforth, we denote the operator $\mathbb{E}^*[\cdot]$ by $\mathbb{E}[\cdot]$. Let $R_{d,p}$ denote an independent copy of $Q_{d,p}$. Then $\alpha_{i,d}^2=\alpha_{1,d}^2\mathbb{E}\left[(Q_{d,p}R_{d,p})^{i-1}\right]$. Using the identity $\sum_{i=1}^{k}(-1)^{i-1}\binom{k}{i}x^{i-1}=\frac{1-(1-x)^k}{x}$, we translate the inequality from (4.11) via the identity (4.14) into

$$\liminf_{d\to\infty}\inf_{p\in[0,1]}\ \alpha_{1,d}\mathbb{E}\left[\frac{1-(1-Q_{d,p})^k}{Q_{d,p}}\right]-\frac{\alpha_{1,d}^2}{2}\mathbb{E}\left[\frac{1-(1-Q_{d,p}R_{d,p})^k}{Q_{d,p}R_{d,p}}\right]\geq 0\,. \tag{4.15}$$

Now we make a choice for the value of the coupling parameter p for each value of d. Fix a parameter $u>0$ and for all large d, pick $p=p(d,u)$ in the construction of the coupling so that

$$\mathbb{E}\left[Q_{d,p}^u\right]=1/\alpha.$$

This can be done due to the continuity of $p\to\mathbb{E}\left[Q_{d,p}^u\right]$ as given by Lemma 4.4 and the assumption that $\alpha>1$. We denote $Q_{d,p(d,u)}$ by Q_d for convenience. As of now it is not at all clear as to why we set p in this manner. However, in the following argument we will see that setting p this way is a judicious choice.

We now explain how to take a limit of Q_d as $d\to\infty$ so that we may analyze (4.15) in terms of k. Since probability distributions on $[0,1]$ are compact in the weak topology by Prokhorov's Theorem, we choose a subsequence (Q_{d_i},R_{d_i}) such that it converges in the weak limit to (Q,R). The random variables Q and R are independent and identically distributed with values in $[0,1]$.

Set

$$s_k(x) = \frac{1-(1-x)^k}{x} \text{ for } 0 \le x \le 1.$$

Observe that $s_k(x) = 1+(1-x)+\cdots+(1-x)^{k-1}$. Thus, $s_k(x)$ is continuous, decreasing on $[0,1]$ with maximum value $s_k(0) = k$ and minimum value $s_k(1) = 1$. Continuity of s_k implies that $\mathbb{E}[s_k(Q_{d_i})] \to \mathbb{E}[s_k(Q)]$ and $\mathbb{E}[s_k(Q_{d_i}R_{d_i})] \to \mathbb{E}[s_k(QR)]$ by the definition of weak convergence of random variables. We now take limits along the subsequence d_i in (4.15). The quantity $\alpha_{1,d}$ may not have a limit but since we assume, for sake of a contradiction, that $\alpha_{1,d} \ge \alpha > 1$, we deduce from taking limits in the d_i that for all $k \ge 1$

$$\mathbb{E}[s_k(Q)] \ge \frac{\alpha}{2}\mathbb{E}[s_k(QR)]. \tag{4.16}$$

4.4. **Final step: large k analysis.** We want to take the limit in k of (4.16) as well but we must be careful. The function $s_k(x)$ monotonically converges to $1/x$ for $x \in [0,1]$. So by taking limits in k in (4.16) and using the independence of Q and R we conclude that $2\mathbb{E}[1/Q] \ge \alpha\mathbb{E}[1/Q]^2$. Of course, we do not know a priori that $\mathbb{E}[1/Q] < \infty$. Even if it were, we can only conclude that $\alpha \le 2/\mathbb{E}[1/Q]$, which leaves us with the seemingly contradictory task of showing that $\mathbb{E}[1/Q]$ is large but finite.

To deal with these issues we have to use the fact that we have chosen $p = p(d,u)$ such that $\mathbb{E}[Q_d^u] = 1/\alpha$. In particular, as $x \to x^u$ is continuous for $0 \le x \le 1$, we see that $\mathbb{E}[Q^u] = \lim_{i\to\infty}\mathbb{E}\left[Q_{d_i}^u\right] = 1/\alpha$. We have to control the expectation $\mathbb{E}[1/Q]$ through our control of $\mathbb{E}[Q^u]$.

Three cases can arise: $\mathbb{P}[Q=0] > 0$, or $\mathbb{P}[Q=0] = 0$ but $\mathbb{E}[1/Q] = \infty$, or $\mathbb{E}[1/Q] < \infty$.

Case 1: $\mathbb{P}[Q=0] = q > 0$. The majority of the contribution to $\mathbb{E}[s_k(Q)]$ results from $\{Q=0\}$. More precisely, $\frac{s_k(x)}{k} \to \mathbf{1}_{x=0}$ as $k \to \infty$, and $\frac{s_k(x)}{k} \in [0,1]$ for all k and $x \in [0,1]$. From the bounded convergence theorem we deduce that $\mathbb{E}[s_k(Q)/k] \to \mathbb{P}[Q=0]$ and $\mathbb{E}[s_k(QR)/k] \to \mathbb{P}[QR=0]$ as $k \to \infty$. The latter probability is $2q - q^2$ due to Q and R being independent and identically distributed. By dividing the inequality in (4.16) by k and taking a limit we conclude that

$$2q - \alpha(2q-q^2) \ge 0, \text{ or equivalently, } \alpha \le \frac{2}{2-q}.$$

Now, since for $x \in [0,1]$, we have that $\mathbf{1}_{x=0} \le 1 - x^u$. It follows from here that $q \le 1 - \mathbb{E}[Q^u] = 1 - 1/\alpha$. Hence,

$$\alpha \le \frac{2}{2-q} \le \frac{2}{1+\alpha^{-1}}.$$

Simplifying the latter inequality gives $\alpha \le 1$; a contradiction.

Case 2: $\mathbb{P}[Q=0] = 0$ but $\mathbb{E}\left[\frac{1}{Q}\right] = \infty$. Now, most of the contribution to $\mathbb{E}[s_k(Q)]$ occurs when Q is small. Recall that $s_k(x) \nearrow 1/x$ as $k \to \infty$. By the monotone convergence theorem, $\mathbb{E}[s_k(Q)] \to \infty$ as $k \to \infty$.

Fix $0 < \epsilon < 1$, and write $s_k(x) = s_{k,\le\epsilon}(x) + s_{k,>\epsilon}(x)$ where $s_{k,\le\epsilon}(x) = s_k(x)\,\mathbf{1}_{x\le\epsilon}$. Note that $s_{k,>\epsilon}(x) \le \epsilon^{-1}$ for all k. We have that

$$\mathbb{E}[s_k(Q)] = \mathbb{E}[s_{k,\le\epsilon}(Q)] + \mathbb{E}[s_{k,>\epsilon}(Q)] \le \mathbb{E}[s_{k,\le\epsilon}(Q)] + \epsilon^{-1}. \tag{4.17}$$

Therefore, $\mathbb{E}[s_{k,\le\epsilon}(Q)] \to \infty$ with k since $\mathbb{E}[s_k(Q)] \to \infty$.

Observe from the positivity of s_k that

$$\mathbb{E}[s_k(QR)] \ge \mathbb{E}[s_k(QR); Q \le \epsilon, R > \epsilon] + \mathbb{E}[s_k(QR); Q > \epsilon, R \le \epsilon].$$

The latter two terms are equal by symmetry, so $\mathbb{E}[s_k(QR)] \ge 2\mathbb{E}[s_k(QR); Q \le \epsilon, R > \epsilon]$. The fact that $s_k(x)$ is decreasing in x and $R \le 1$ imply that $s_k(QR) \ge s_k(Q)$. Together with the independence of Q and R we deduce that

$$\mathbb{E}[s_k(QR); Q \le \epsilon, R > \epsilon] \ge \mathbb{E}[s_k(Q); Q \le \epsilon, R > \epsilon] = \mathbb{E}[s_{k,\le\epsilon}(Q)]\,\mathbb{P}[R > \epsilon].$$

Consequently,

$$\mathbb{E}\left[s_k(QR)\right] \geq 2\mathbb{E}\left[s_{k,\leq\epsilon}(Q)\right]\mathbb{P}\left[Q > \epsilon\right]. \tag{4.18}$$

The inequality in (4.16) is $\frac{\alpha}{2} \leq \frac{\mathbb{E}[s_k(Q)]}{\mathbb{E}[s_k(QR)]}$. The bounds from (4.17) and (4.18) imply that

$$\frac{\alpha}{2} \leq \frac{\mathbb{E}\left[s_{k,\leq\epsilon}(Q)\right] + \epsilon^{-1}}{2\mathbb{E}\left[s_{k,\leq\epsilon}(Q)\right]\mathbb{P}\left[Q > \epsilon\right]}.$$

Since $\mathbb{E}\left[s_{k,\leq\epsilon}(Q)\right] \to \infty$ with k we can take a limit in k to conclude that

$$\alpha \leq \frac{1}{\mathbb{P}\left[Q > \epsilon\right]}.$$

As $\epsilon \to 0$ the probability $\mathbb{P}\left[Q > \epsilon\right] \to \mathbb{P}\left[Q > 0\right] = 1$, by assumption. Thus, $\alpha \leq 1$, a contradiction.

Case 3: $\mathbb{E}\left[1/Q\right]$ is finite. Since $s_k(x)$ monotonically converges to $1/x$ as $k \to \infty$, we deduce from (4.16) that

$$2\mathbb{E}\left[\frac{1}{Q}\right] - \alpha\mathbb{E}\left[\frac{1}{QR}\right] \geq 0.$$

Since $\mathbb{E}\left[1/(QR)\right] = \mathbb{E}\left[1/Q\right]^2$, the inequality above becomes $\alpha \leq 2\mathbb{E}\left[1/Q\right]^{-1}$. Jensen's inequality implies that $\mathbb{E}\left[1/Q\right]^{-1} \leq \mathbb{E}\left[Q^u\right]^{1/u}$. However, $\mathbb{E}\left[Q^u\right] = 1/\alpha$ and this means that $\alpha \leq 2^{\frac{u}{u+1}}$. As the previous two cases led to contradictions we conclude that $\alpha \leq 2^{\frac{u}{u+1}}$ for all $u > 0$. However, as $u \to 0$ we see that $\alpha \leq 1$. This final contradiction completes the proof of Theorem 1.1.

Proof of Corollary 1.2. If $\mathbf{Y}$ is an FIID percolation process on $\mathbb{T}_d$ that has finite components with probability one, then the components of $\mathbf{Y}$ are all finite trees. Thus, the component of the root provides a measure on finite, rooted trees with the root being picked uniformly at random (due to invariance and transitivity). If a tree has n vertices then the expected degree of a uniform random root is $2(n-1)/n \leq 2$. This implies that the average degree of the root of $\mathbb{T}_d$, given that it is included in $\mathbf{Y}$, cannot be larger than 2. As such, the hypothesis of Theorem 1.1 applies in this situation.

4.5. **Density bounds for small values of d.** We conclude this section with remarks on FIID percolation processes on $\mathbb{T}_d$ for 'small' values of d. Although the entropy inequality may be used in principal for every d to derive bounds on the density of FIID percolation on $\mathbb{T}_d$ with finite clusters, these upper bounds are not the best available for small values of d. The best known explicit upper bounds for small values of d ($d \leq 10$) is due to Bau et. al. [6]. They are derived by using more exhaustive counting arguments on random d-regular graphs than those which provide our entropy inequality. The cases for $d = 3$ or $d = 4$ are rather interesting. The upper bound of $2/3$ for $\mathbb{T}_4$ have not be ruled out by further counting arguments, but the best known lower bound is 0.6045 due to a construction of Hoppen and Wormald [17]. We suspect that the upper bound of $2/3$ is optimal.

For $d = 3$, Bau et. al. provide a upper bound of $3/4$. Endre Csóka has confirmed the upper bound of $3/4$ as optimal in personal communication . Csóka's construction may be derived from an appropriate edge orientations of $\mathbb{T}_3$. There exists, as a weak limit of FIID processes, an invariant orientation of the edges of $\mathbb{T}_3$ such that the out-degree of every vertex is either 1 or 3. The percolation process consists of the vertices with out-degree 1. Invariance implies that their density is $3/4$. The clusters are infinite but have at most two topological ends (which means that no cluster contains three disjoint infinite rays emanating from a common vertex). The clusters can then be made finite by performing a Bernoulli percolation on the vertices of out-degree 1 at any density less than 1. The density of the resulting percolation process is arbitrarily close to $3/4$.

Although Csóka's construction realizes the optimal percolation process as a weak limit of FIID processes, it is not necessarily true that the limiting process is itself an FIID process. There exist invariant, but not FIID, processes on $\mathbb{T}_d$ that are weak limits of FIID process [15, Theorem 6].

5. FIID edge orientations of $\mathbb{T}_d$

Following Csóka's construction of percolation processes via edge orientations, we may ask what other types of FIID edge orientations are possible on $\mathbb{T}_d$. In this section we will show that $\mathbb{T}_d$ admits an FIID edge orientation with no sources or sinks.

FIID processes on the edges on $\mathbb{T}_d$ are defined analogous to vertex-indexed processes. Since $\mathrm{Aut}(\mathbb{T}_d)$ acts transitively on the edges of $\mathbb{T}_d$ the definition for vertex-indexed processes carries over naturally. An edge indexed FIID process $\Phi \in \chi^{E(\mathbb{T}_d)}$ is thus determined by a measurable function $F : [0,1]^{E(\mathbb{T}_d)} \to \chi^{E(\mathbb{T}_d)}$ satisfying $F(\gamma \cdot \omega) = \gamma \cdot F(\omega)$ for all $\gamma \in \mathrm{Aut}(\mathbb{T}_d)$. Φ is then defined as $\Phi = F(\mathbf{X})$, where $\mathbf{X}$ is a random labelling of the edges of $\mathbb{T}_d$.

Given any edge orientation of a graph G, a vertex v is called a source if all edges that are incident to v are oriented away from it. If all incident edges are oriented towards v, then v is called a sink. For $d \geq 3$, we use the existence of FIID perfect matchings [22] to construct an FIID edge orientation of $\mathbb{T}_d$ with no sources or sinks. Such an orientation cannot exists for $\mathbb{T}_2$. Indeed, $\mathbb{T}_2$ has only two edge orientations with no sources or sinks: either all edges point to the right or all point to the left. However, the uniform distribution on these two orientations cannot be an FIID process because it is not mixing.

Theorem 5.1. *For every $d \geq 3$ there exists an FIID orientation of the edges of $\mathbb{T}_d$ that contains no sources or sinks.*

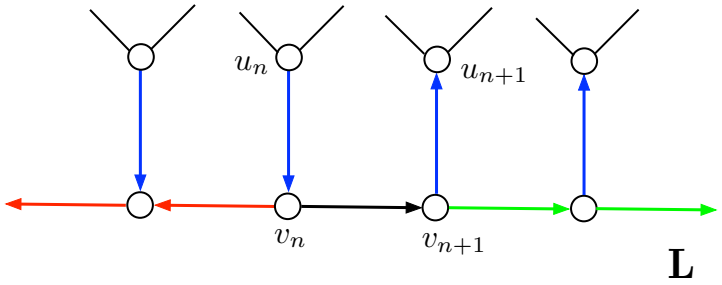

FIGURE 1. Orientation of a bi-infinite path $\mathbf{L}$. Edges of the perfect matching that meet $\mathbf{L}$ are in blue. The edges of L_n are red and edges of L_{n+1} are green. The black edge between v_n and v_{n+1} connects L_n to L_{n+1}.

Proof. Lyons and Nazarov prove that $\mathbb{T}_d$ has an FIID perfect matching [22]. We consider an FIID perfect matching of $\mathbb{T}_d$ and orient its edges independently and uniformly at random. Then we remove the oriented edges of the matching and are left with a forest whose components are isomorphic to $\mathbb{T}_{d-1}$. We continue with this procedure of selecting an FIID perfect matching from the components, orienting those edges at random and then removing them from the graph, until we are left with a forest whose components are isomorphic to $\mathbb{T}_3$. Thus, we have to orient the edges of $\mathbb{T}_3$ with no sources or sinks as an FIID process to finish.

Henceforth, we explain how to orient the edges of $\mathbb{T}_3$ in such a manner. Begin by finding an FIID perfect matching on $\mathbb{T}_3$ and then orient the edges of the matching independently and uniformly at random. The un-oriented edges span a 2-factor, that is, its components are un-oriented bi-infinite paths. Consider any

such path $\mathbf{L}$. The oriented perfect matching partitions $\mathbf{L}$ into contiguous finite paths $L_n, n \in \mathbb{Z}$, which are characterized by the following properties.

First, L_n is incident to L_{n+1} via an edge $\{v_n, v_{n+1}\}$ in $\mathbf{L}$. Second, for every n, the edges of the perfect matching that meet L_n are all oriented in 'parallel', either all pointing towards L_n or away from it. Finally, the edges of the matching that meet L_{n+1} are oriented in the opposite direction compared to the (common) orientation of those edges of the matching that meet L_n. This setup is shown in Figure 1. The reason the paths L_n are finite is because the edges of the matching that are incident to $\mathbf{L}$ are oriented independently. We can think of the L_ns that are pointing towards the matching as the clusters of a Bernoulli percolation on $\mathbf{L}$ at density 1/2.

To complete the orientation we first orient all the edges on the path L_n in the same direction. As there are two possible directions, we choose one at random. This is done independently for each L_n on every un-oriented bi-infinite path $\mathbf{L}$. The finiteness of the L_n is crucial to ensure that these orientations can be done as an FIID. Following this, any vertex in the interior of the path L_n has one edge directed towards it and another directed away. This ensures that such a vertex cannot form a source or a sink.

The vertices that can form sources or sinks are those at the endpoints of any path L_n. Also, the un-oriented edges that remain are the $\{v_n, v_{n+1}\}$ which connect two contiguous paths L_n and L_{n+1} on a common bi-infinite path $\mathbf{L}$. Observe that the endpoints of these edges are precisely the vertices that can form sources and sinks (see the black edge in Figure 1). Given such an edge $\{v_n, v_{n+1}\}$ let u_n (resp. u_{n+1}) be the neighbour of v_n (resp. v_{n+1}) such that the edge $\{v_n, u_n\}$ (resp. $\{v_{n+1}, u_{n+1}\}$) lies in the matching. The edges $\{v_n, u_n\}$ and $\{v_{n+1}, u_{n+1}\}$ have opposite orientations by design (see Figure 1). Suppose that $u_n \to v_n$ and $v_{n+1} \to u_{n+1}$. We orient $\{v_n, v_{n+1}\}$ as $v_n \to v_{n+1}$, thus ensuring that both v_n and v_{n+1} have one incoming edge and another outgoing one. This completes the orientation of $\mathbb{T}_3$ and produces no sources or sinks.

□

Acknowledgements. The Authors research has been supported by an NSERC CGS grant. The author is grateful to Endre Csóka and Viktor Harangi for explaining the construction of the optimal percolation process of density 3/4 in Section 4.5. The author also thanks Bálint Virág for helpful discussions and the anonymous referees for helpful suggestions regarding the exposition.

References

[1] M. Bayati, D. Gamarnik, and P. Tetali, Combinatorial approach to the interpolation method and scaling limits in sparse random graphs, *Ann. Probab.* **41** (2013), pp. 4080–4115.

[2] B. Bollobás, Random graphs, 2nd ed., Cambridge University Press, 2001.

[3] B. Bollobás, The independence ratio of regular graphs, *Proc. Amer. Math. Soc.* **83** no. 2 (1981), pp. 433–436.

[4] Á. Backhausz, B. Szegedy, and B. Virág, Ramanujan graphings and correlation decay in local algorithms, *Random Struct. Alg.* (2014), doi: 10.1002/rsa.20562.

[5] Á. Backhausz and B. Szegedy, On large girth graphs and random processes on trees, preprint (2014), arXiv:1406.4420.

[6] S. Bau, N.C. Wormald, and S. Zhou, Decycling number of random regular graphs, *Random Struct. Alg.* **21** (2002), pp. 397–413.

[7] L. P. Bowen, A measure-conjugacy invariant for free group actions, *Ann. of Math.* **171**(2) (2010), pp. 1387–1400.

[8] E. Csóka, B. Gerencsér, V. Harangi, and B. Virág, Invariant Gaussian processes and independent sets on regular graphs of large girth, *Random Struct. Alg.* (2014), doi: 10.1002/rsa.20547.

[9] E. Csóka and G. Lippner, Invariant random matchings in Cayley graphs, preprint (2012), arXiv:1211.2374.

[10] G. Elek and G. Lippner, Borel oracles. An analytical approach to constant-time algorithms, *Proc. Amer. Math. Soc.* **138** (2010), pp. 2939–2947.

[11] A.M. Frieze and T. Łuczak, On the independence and chromatic numbers of random regular graphs, *J. Combin. Theory Ser. B* **54** (1992), pp. 123–132.

[12] D. Gamarnik and D.A. Goldberg, Randomized greedy algorithms for independent sets and matchings in regular graphs: Exact results and finite girth corrections, *Combin. Probab. Comput.* **19** (2010), pp. 61–85.

[13] D. Gamarnik and M. Sudan, Limits of local algorithms over sparse random graphs, *Proc. 5th Innovations in Theoretical Computer Science Conf.*, ACM (2014) pp. 369–376, arXiv:1304.1831.

[14] O. Häggström, Infinite clusters in dependent automorphism invariant percolation on trees, *Annals of Probability* **25**(3) (1997), pp. 1423–1436.

[15] V. Harangi and B. Virág, Independence ratio and random eigenvectors in transitive graphs, preprint (2013) arXiv:1308.5173

[16] H. Hatami, L. Lovász, and B. Szegedy, Limits of local-global convergent graph sequences, *Geom. Funct. Anal.* **24**(1) (2014), pp. 269–296.

[17] C. Hoppen and N. Wormald, Induced forests in regular graphs with large girth, *Combin. Probab. Comput.* **17**(3) (2008), pp. 389–410.

[18] C. Hoppen and N. Wormald, Local algorithms, regular graphs of large girth, and random regular graphs, preprint (2013) arXiv:1308.0266v2.

[19] S. Janson, T. Łuczak, and A. Rucinski, Random graphs, John Wiley and Sons, Inc., 2000.

[20] J. Lauer and N. Wormald, Large independent sets in regular graphs of large girth, *J. Combin. Theory Ser. B* **97** (2007), pp. 999–1009.

[21] R. Lyons, Factor of IID on trees, to appear in *Combin. Probab. Comput.* (2015) arXiv:1401.4197.

[22] R. Lyons and F. Nazarov, Perfect matchings as iid factors on non-amenable groups, *European J. Combin.* **32** (2011), pp. 1115–1125.

[23] R. Lyons and Y. Peres, Probability on Trees and Networks, Cambridge University Press, 2016, to appear. Available at `http://pages.iu.edu/~rdlyons/`.

[24] M. Rahman, Percolation with small clusters on random graphs, *Graphs and Combinatorics* 32(3) (2016), pp. 1167-1185.

[25] M. Rahman and B. Virág, Local algorithms for independent sets are half–optimal, to appear in *Annals of Probability* (2016) arXiv:1402.0485.

[26] J. B. Shearer, A note on the independence number of triangle free graphs, II, *J. Combin. Theory Ser. B* **53**(2) (1991), pp. 300–307.

[27] M. Talagrand, Mean field models for spin glasses: Volume 1, Basic examples, Springer, 2010.

(Mustazee Rahman) Department of Mathematics, University of Toronto, 40 St. George Street, Toronto, ON M5S 2E4, Canada

E-mail address, Mustazee Rahman: `mustazee@math.toronto.edu`